

National Laboratory for Scientific Computing
Graduate Program in Computational Modeling

**Managing large-scale scientific hypotheses as uncertain
and probabilistic data**

By

Bernardo Gonçalves

PETRÓPOLIS, RJ - BRASIL

FEBRUARY - 2015

MANAGING LARGE-SCALE SCIENTIFIC HYPOTHESES AS
UNCERTAIN AND PROBABILISTIC DATA

Bernardo Gonçalves

THESIS SUBMITTED TO THE EXAMINING COMMITTEE IN PARTIAL
FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF DOC-
TOR OF SCIENCES IN COMPUTATIONAL MODELING.

Approved by:

Prof. Fabio Porto, D.Sc.
(Chair)

Prof. Pedro L. Dias, Ph.D.

Prof. Marco A. Casanova, Ph.D.

Prof. Ana Carolina Salgado, D.Sc.

PETRÓPOLIS, RJ - BRASIL
FEBRUARY - 2015

© 2015, Bernardo Nunes Gonçalves.

All rights reserved.

Gonçalves, Bernardo

G635m Managing large-scale scientific hypotheses as uncertain and probabilistic data / Bernardo Gonçalves. Petropolis, RJ. : National Laboratory for Scientific Computing, 2015.

xvii, 128p. : il.; 29 cm

Orientador: Fabio Porto

Thesis (D.Sc.) – National Laboratory for Scientific Computing, 2015.

1. Hypothesis management. 2. Predictive analytics. 3. Uncertain and probabilistic data. 4. Causal reasoning. 5. Probabilistic database design. I. Porto, Fabio. II. LNCC/MCT. III. Título.

CDD – 629.8

Dedictory

To my parents Tania and Francisco,
and to my special love, Marcelle,
for being an island of certainty in an
uncertain world.

Acknowledgments

This thesis work has been supported by LNCC (Graduate Program in Computational Modeling), CNPq (grant 141838/2011-6), FAPERJ (grant ‘Nota 10’ E-26/100.286/2013) and IBM (Ph.D. Fellowship 2013/2014).

I would like to express my gratitude to my advisor Fabio Porto for the gift of the challenging topic of this thesis, so special to me. I am grateful to him for being an inspiring advisor, and for nicely influencing me towards database research. I also thank my thesis committee for their attention and time devoted in the assessment of my work. I am indebted to Ana Maria Moura for her generous advice and support throughout my Ph.D. research, and to Frederico C. Silva and Adolfo Simões for their support to my research project at DEXL/LNCC.

I would like to thank very much all researchers at LNCC who give lectures in the graduate program, specially Prof. José Karam Filho for teaching me generously about the roots of mathematical modeling. They have contributed significantly to my education as a cross-disciplinary thinker and the shaping of my scientific and mathematical skills. I thank my Ph.D. colleagues and friends at LNCC, specially Eduardo Lima, Ramon Costa, Klaus Wehmuth, Raquel Lopes, Karine Guimarães and Diego Paredes for their companionship and joy shared in the pursuit of their Ph.D. theses. I would also like to gratefully recall my earlier professors at UFES, José Gonçalves, Giancarlo Guizzardi, Rosane Caruso and Berilhes Garcia for shaping the most essential building blocks in my education.

Finally, I would like to thank my family for their support: my father Francisco, my example of simplicity and goodness; my mother, Tania, for her tenacity and true love; Leo and Sche, my dear brother and sister, true union for life; my grandmothers Zeca, Orizontina and Rosita for their love and prayers. I thank also Marcel and Maris for the greatest gift, their daughter and my near future wife, Marcelle – the best partner one could ever hope for.

“Originally, there was just experimental science, and then there was theoretical science, with Kepler’s Laws, Newton’s Laws of Motion, Maxwell’s equations, and so on. Then, for many problems, the theoretical models grew too complicated to solve analytically, and people had to start simulating. These simulations have carried us through much of the last half of the last century. At this point, these simulations are generating a whole lot of data, along with a huge increase in data from the experimental sciences.”

— Jim Gray, 2007

Abstract of Thesis presented to LNCC/MCT in partial fulfillment of the requirements for the degree of Doctor of Sciences (D.Sc.)

MANAGING LARGE-SCALE SCIENTIFIC HYPOTHESES AS UNCERTAIN AND PROBABILISTIC DATA

Bernardo Gonçalves

February - 2015

Advisor: Fabio Porto, D.Sc.

In view of the paradigm shift that makes science ever more data-driven, in this thesis we propose a synthesis method for encoding and managing large-scale deterministic scientific hypotheses as uncertain and probabilistic data.

In the form of mathematical equations, hypotheses symmetrically relate aspects of the studied phenomena. For computing predictions, however, deterministic hypotheses can be abstracted as functions. We build upon Simon's notion of structural equations in order to efficiently extract the (so-called) causal ordering between variables, implicit in a hypothesis structure (set of mathematical equations).

We show how to process the hypothesis predictive structure effectively through original algorithms for encoding it into a set of functional dependencies (fd's) and then performing causal reasoning in terms of acyclic pseudo-transitive reasoning over fd's. Such reasoning reveals important causal dependencies implicit in the hypothesis predictive data and guide our synthesis of a probabilistic database. Like in the field of graphical models in AI, such a probabilistic database should be normalized so that the uncertainty arisen from competing hypotheses is decomposed into factors and propagated properly onto predictive data by recovering its joint probability distribution through a lossless join. That is motivated as a design-theoretic principle for data-driven hypothesis management and predictive analytics.

The method is applicable to both quantitative and qualitative deterministic hypotheses and demonstrated in realistic use cases from computational science.

Resumo da Tese apresentada ao LNCC/MCT como parte dos requisitos necessários para a obtenção do grau de Doutor em Ciências (D.Sc.)

GERÊNCIA DE HIPÓTESES CIENTÍFICAS DE LARGA-ESCALA COMO DADOS INCERTOS E PROBABILÍSTICOS

Bernardo Gonçalves

Fevereiro, 2015

Orientador: Fabio Porto, D.Sc.

Tendo em vista a mudança de paradigma que faz da ciência cada vez mais guiada por dados, nesta tese propomos um método para codificação e gerência de hipóteses científicas determinísticas de larga escala como dados incertos e probabilísticos.

Na forma de equações matemáticas, hipóteses relacionam simetricamente aspectos do fenômeno de estudo. Para computação de predições, no entanto, hipóteses determinísticas podem ser abstraídas como funções. Levamos adiante a noção de Simon de equações estruturais para extrair de forma eficiente a então chamada ordenação causal implícita na estrutura de uma hipótese.

Mostramos como processar a estrutura preditiva de uma hipótese através de algoritmos originais para sua codificação como um conjunto de dependências funcionais (df's) e então realizamos inferência causal em termos de raciocínio acíclico pseudo-transitivo sobre df's. Tal raciocínio revela importantes dependências causais implícitas nos dados preditivos da hipótese, que conduzem nossa síntese do banco de dados probabilístico. Como na área de modelos gráficos (IA), o banco de dados probabilístico deve ser normalizado de tal forma que a incerteza oriunda de hipóteses alternativas seja decomposta em fatores e propagada propriamente recuperando sua distribuição de probabilidade conjunta via junção 'lossless.' Isso é motivado como um princípio teórico de projeto para gerência e análise de hipóteses.

O método proposto é aplicável a hipóteses determinísticas quantitativas e qualitativas e é demonstrado em casos realísticos de ciência computacional.

Table of Contents

1	Introduction	1
1.1	Problem Space and Specific Goals	4
1.2	Thesis Statement	12
1.3	Thesis Contributions	13
1.4	Thesis Outline	14
2	Vision: Hypotheses as Data	15
2.1	Running Example	18
2.2	Hypothesis Encoding	20
2.3	Reasoning over FD's	21
2.4	Uncertainty Introduction	22
2.5	Predictive Analytics	25
2.6	Related Work	27
2.7	Summary: Key Points	30
3	Hypothesis Encoding	32
3.1	Preliminaries: Structural Equations	32
3.2	The Problem of Causal Ordering	36
3.3	Total Causal Mappings	39
3.4	The Encoding Scheme	43
3.5	Experiments	45
3.6	Related Work	46
3.7	Summary of Results	48
4	Causal Reasoning over FD's	49
4.1	Preliminaries: Armstrong's Inference Rules	49

4.2	Acyclic Pseudo-Transitive Reasoning	50
4.3	Equivalence with Causal Ordering	55
4.4	Experiments	58
4.5	Related Work	58
4.6	Summary of Results	60
5	Probabilistic Database Synthesis	61
5.1	Preliminaries: U-Relations and Probabilistic WSA	61
5.2	Running Example	63
5.3	U-Factorization	66
5.4	U-Propagation	70
5.5	Design-Theoretic Properties	73
5.6	Related Work	76
5.7	Summary of Results	78
6	Applicability	79
6.1	The Physiome Project as a Testbed	79
6.2	Case Studies	80
6.3	System Prototype	89
6.4	Experiments	95
6.5	Discussion	98
6.6	Conclusions	101
7	Conclusions	102
7.1	Revisiting the Research Questions	102
7.2	Significance and Limitations	105
7.3	Open Problems and Future Work	106
7.4	Final Considerations	107

Bibliography	108
--------------	-----

Appendix

A Detailed Proofs	116
A.1 Proofs of Hypothesis Encoding	116
A.2 Proofs of Causal Reasoning	119
A.3 Proofs of Probabilistic DB Synthesis	126

List of Figures

	Figure
1.1 Multi-fold view of a deterministic scientific hypothesis.	2
1.2 A view of the scientific method life cycle	5
1.3 The usual data ingesture pipeline of simulation data management .	7
1.4 Design-theoretic pipeline for processing hypotheses as uncertain and probabilistic data	8
2.1 Scientific hypotheses seen as alternative functions to predict data. .	16
2.2 Predictive analytics in a data-intensive hypothesis evaluation study.	18
2.3 Descriptive (textual) data of Example 1.	19
2.4 ‘Big’ fact table H_1 loaded with simulation raw data	19
2.5 ‘Explanation’ relational table H_0	23
2.6 Result set of query Q_1 on simulation trial dataset for hypothesis H_1 .	24
2.7 U-relational predictive tables rendered by query using the fd’s. . . .	25
2.8 Analytics on predicted position s conditioned on observation.	26
3.1 “Directed causal graphs” associated with the two systems.	33
3.2 Running Simon’s Causal Ordering Algorithm (COA)	35
3.3 Directed causal graph G_φ induced by mapping φ for structure \mathcal{S} . .	36
3.4 Bipartite graph G of structure \mathcal{S} from Example 3.	37
3.5 Another hypothesis structure example.	39
3.6 Complete matching M for structure \mathcal{S} from Example 3.	42
3.7 Encoded fd set Σ (cf. Alg. 3) for the structure from Example 3. . . .	44
3.8 Performance of hypothesis encoding (in logscale).	46
4.1 Fd set Σ encoding (cf. Alg. 3) the structure of Fig. 3.2a and its folding $\Sigma^{\text{q}\rightarrow}$ derived by Alg. 5.	52

4.2	Fd set Σ encoding the structure of Fig. 3.2a and the folding $\Upsilon(\Sigma)^{\text{q}}$ of its v -projection.	57
4.3	Performance of acyclic causal reasoning over fd's (logscale).	59
5.1	U-relation generated by the repair-key operation.	63
5.2	Fd sets encoded from the given structures $\mathcal{S}_k(\mathcal{E}_k, \mathcal{V}_k)$ for hypotheses $k = 1..3$ from Example 6.	65
5.3	'Big' fact table H_3 of hypothesis $k=3$ from Example 6 loaded with trial datasets identified by special attribute tid	65
5.4	'Big' fact table H_3 of hypothesis $k=3$ from Example 6 with u-factors $\{b, d\}$ and $\{p, r\}$ emphasized (resp.) in colors green and red.	67
5.5	Fd set Ω_3 (compare with Σ_3) and its folding Ω_3^{q}	68
5.6	U-factor projections rendered for hypothesis $v = 3$	69
5.7	U-relations rendered for hypothesis $v = 3$	72
6.1	Plot of hemoglobin oxygen saturation hypotheses	81
6.2	Descriptive (textual) data of Example 8.	81
6.3	Fd set Σ_{28} of hypothesis $v=28$	82
6.4	Fd set Σ_{31} of hypothesis $v=31$	82
6.5	Fd set Σ_{32} of hypothesis $v=32$	82
6.6	Result set of hypothesis management query Q1.	83
6.7	Results of analytical study on the hemoglobin phenomenon.	83
6.8	Plot of baroreflex hypothesis for Dahl SS Rat	84
6.9	Descriptive (textual) data of Example 9.	84
6.10	Result set of hypothesis management query Q2.	85
6.11	Results of analytical study on the baroreflex phenomenon.	85
6.12	Fd set Σ_{1001} of hypothesis $v=1001$	86
6.13	Plot of myogenic behavior hypothesis	87
6.14	Descriptive (textual) data of Example 10.	87
6.15	Fd set Σ_{60} of hypothesis $v=60$	88
6.16	Fd set Σ_{89} of hypothesis $v=89$	88
6.17	Result set of hypothesis management query Q3.	89
6.18	Results of analytics on the vessel's myogenic behavior phenomenon.	89

6.19	Lynx-Hare population observed	91
6.20	Lynx-Hare population observed	91
6.21	Screenshots of this first prototype of the Υ -DB system.	92
6.22	Descriptive (textual) data of Example 10.	93
6.23	Fd set Σ_1 of hypothesis $v=1$	93
6.24	Fd set Σ_2 of hypothesis $v=2$	93
6.25	Fd set Σ_3 of hypothesis $v=3$	94
6.26	Results of analytics on the US population phenomenon.	95
6.27	Results of analytics on the Hudson's Bay lynx population phenomenon	95
6.28	Performance behavior of Υ -DB on a Physiome testbed scenario. . .	96
6.29	Physiome hypotheses used in the experiments.	97
6.30	Example of Boolean Network hypothesis.	100
6.31	Example of Boolean Network model	101

List of Tables

Table

1.1	Simulation data management vs. hypothesis data management.	. . .	5
-----	--	-------	---

Acronyms

- fd: functional dependency
- p-DB: probabilistic database
- p-WSA: probabilistic world set algebra
- MayBMS: U-relational database management system
- AI: Artificial Intelligence
- GM: graphical models (e.g., Bayesian Networks)
- ETL: extract, transform, load
- OLAP: On-Line Analytical Processing
- DW: Data Warehouse
- SEM: structural equation model
- COA: causal ordering algorithm
- \mathcal{S}_k : structure of hypothesis k
- \mathcal{E} : set of equations in a structure
- \mathcal{V} : set of variables appearing in the equations in a structure
- $Vars(f)$: set of variables appearing in equation f
- $\varphi: \mathcal{E} \rightarrow \mathcal{V}$: total causal mapping from equations to variables in a structure
- C_φ : set of causal dependencies
- G_φ : causal graph induced by φ
- H_k : ‘big’ fact table of hypothesis k
- \mathbf{H} : set of hypothesis ‘big’ fact tables
- Y_k^ℓ : U-relation synthesized for hypothesis k
- \mathbf{Y}_k : set of U-relations synthesized for hypothesis k
- H_0 : relational ‘explanation’ table

- Y_0 : U-relational ‘explanation’ table
- ϕ : phenomenon identifier
- v : hypothesis identifier
- $\Sigma, \Gamma, \Delta, \Omega$: fd sets
- $\Sigma^>$: subset of the closure of an fd set derived by pseudo-transitivity
- Σ^+ : closure of an fd set
- Σ^{tr} : the folding of an fd set
- synthesis ‘4U’: synthesis for uncertainty
- tid: hypothesis trial identifier
- BCNF: Boyce-Codd Normal Form
- MathML: Mathematical Markup Language
- TCM: total causal mapping algorithm
- lhs, rhs: left-hand side, right-hand side
- $\text{sch}(R)$: data columns of table R
- $\overline{V_i D_i}$: condition columns of a table

Chapter 1

Introduction

In view of the paradigm shift that makes science ever more data-driven [1], *in this thesis we demonstrate that large deterministic scientific hypotheses can be effectively encoded and managed as a kind of uncertain and probabilistic data.*

Deterministic hypotheses can be formed as principles or ideas, then expressed mathematically and implemented in a program that is run to give their *decisive* form of data (see Fig. 1.1). Hypotheses can also be learned in large scale, as exhibited in the **Eureqa** project [2]. Examples of ‘structured deterministic hypotheses’ include tentative mathematical models in physics, engineering and economical sciences, or conjectured boolean networks in molecular biology and social sciences. These are important reasoning devices, as they are solved to generate valuable predictive data for decision making in science and increasingly in business as well.

In fact, we can refer nowadays to a broad, modern context of *data science* [3] and *big data* [4] in which the complexity and scale of so-called ‘data-driven’ problems require proper data management tools for the predicted data to be analyzed effectively. In this thesis, we pay attention to a quite general class of (tentative) *computational science* models,¹ and we look at them in an original way as a distinguished kind of data source.

¹ ‘Computational science’ is (sic.) “a rapidly growing multidisciplinary field that uses advanced computing capabilities to understand and solve complex problems” [5]. We may refer to non-stochastic, tentative computational science models throughout this text as ‘structured deterministic hypotheses.’

Law of free fall

“If a body falls from rest, its velocity at any point is proportional to the time it has been falling.”

(i)

$$a(t) = -g$$

$$v(t) = -gt + v_0$$

$$s(t) = -(g/2)t^2 + v_0 t + s_0$$

(ii)

```
for k = 0:n;
    t = k * dt;
    v = -g*t + v_0;
    s = -(g/2)*t^2 + v_0*t + s_0;
    t_plot(k) = t;
    v_plot(k) = v;
    s_plot(k) = s;
end
```

(iii)

FALL	t	v	s
	0	0	5000
	1	-32	4984
	2	-64	4936
	3	-96	4856
	4	-128	4744

(iv)

Figure 1.1. Multi-fold view of a deterministic scientific hypothesis.

It is generally considered that computational science models, interpreted here as hypotheses to explain real-world phenomena, are of strategic relevance [5]. They are usually complex in that they may have hundreds to thousands of intertwined (coupled) variables and be computed along space, time or frequency domains in arbitrarily large scale. It is important to note the distinction between the structure and data levels. Consider, say, Lotka-Volterra’s model, which essentially consists in (Eqs. 1.1) two ordinary differential equations, complemented by seven subsidiary equations $f_1(t)$, $f_2(x_0)$, $f_3(y_0)$, $f_4(b)$, $f_5(p)$, $f_6(r)$, $f_7(d)$ to set the values of its domain variable t and (input) parameters x_0 , y_0 , b , p , r , d .

$$\begin{cases} \dot{x} &= x(b - py) \\ \dot{y} &= y(rx - d) \end{cases} \quad (1.1)$$

In a sense, it can be said fairly simple, as it is characterized by a set \mathcal{E} of equations and a set \mathcal{V} of variables, sized $|\mathcal{E}| = |\mathcal{V}| = 9$. Yet, at the data level this model (cf. Chapter 2) can be made very large just by computing its predictions in a fine time resolution and/or along an extended time window.

As we shall see shortly, the technical challenges associated with this thesis involve (not only but) majorly the structure level where, e.g., such Lotka-Volterra model can be abstracted as a deterministic structure $\mathcal{S}(\mathcal{E}, \mathcal{V})$ with $|\mathcal{S}| = 18$.²

² The structure length $|\mathcal{S}|$ is a measure of how dense the hypothesis structure is, comprising the total sum of the number of variables appearing in each equation.

We are really concerned here with models whose structure \mathcal{S} is in the order of $|\mathcal{S}| \lesssim 1M$, and whose results (data!) shall be difficult to analyze by handicrafted practice. Note that the data level of a model can be set as large as wanted (set the domain resolution and/or extension accordingly), but it shall be necessarily large when its structure is itself large. By ‘large-scale hypotheses’ then we mean tentative deterministic models that are large at structure level.

Overall, such class of hypotheses can be said to qualify to at least four of the five v’s associated to the notion of big data:³ *value*, because of their role in advancing science and technology; *volume*, due to the large scale of modern scientific problems; *variety*, because of their structural heterogeneity, even when they refer to the same phenomena; and *veracity*, due to their uncertainty.

The idea of managing hypotheses ‘as data’ may sound intriguing and in fact it raises a number of research questions of both conceptual and technical nature.⁴ We start by outlining below the conceptual research questions.

- RQ1.** How to define and encode hypotheses ‘as data’? What are the sources of uncertainty that may be present and should be considered?
- RQ2.** How does hypotheses ‘as data’ relate with observational data or, likewise, phenomena ‘as data’ from a database perspective?
- RQ3.** Does every piece of simulated data qualify as a scientific hypothesis? What is the difference between managing ‘simulation’ data from managing ‘hypotheses’ as data?
- RQ4.** Is there available a proper (machine-readable) data format we can use to automatically extract mathematically-expressed hypotheses from?

It has been a challenge of this thesis to provide reasonable answers to these questions, which are brought together into the vision of hypotheses ‘as data’ (we call it the Υ -DB vision) and its use case that we present in Chapter 2, and experiment with in realistic scenarios in Chapter 6.

³ The ‘v’ of *velocity* may appear in connection with machine learning hypotheses, which we discuss in Chapter 6.

⁴ We shall keep record of those questions and revisit them in §7.1.

The Υ -DB vision formulates the problem of hypothesis encoding as a problem of *probabilistic database design*. A number of technical questions arise then.

We introduce now technical context, materials and methods identified and selected in this thesis as a basis to realize the Υ -DB vision in terms of probabilistic database design. We shall outline in the sequel the technical research questions to be answered by the core of the thesis.

1.1. Problem Space and Specific Goals

It has been a goal of this thesis to investigate the capabilities of probabilistic databases to enable hypothesis data management as a particular case of simulation data management. In the sequel, we first characterize the use case of hypothesis data management and then formulate it in terms of probabilistic DB design.

1.1.1 Simulation data management

Simulation laboratories provide scientists and engineers with very large, possibly huge datasets that reconstruct phenomena of interest in high resolution. Notorious examples are the John Hopkins Turbulence Databases [6], and the Human Brain Project (HBP) neuroscience simulation datasets [7]. A core motivation for the delivery of such data is enabling new insights and discoveries through *hypothesis testing against observations*. Nonetheless, while the use case for *exploratory analytics* is currently well understood and many of its challenges have already been coped with so that high-resolution simulation data is increasingly more accessible [8, 9], only very recently, as part of this thesis work, the use case of hypothesis management has been taken into account for *predictive analytics* [10].

In fact, there is a pressing call for innovative technology to integrate (observed) data and (simulated) theories in a unified framework [11, 12, 13]. The point has just been raised by leading neuroscientists in the context of the HBP, who are incisive on the compelling argument that massive simulation databases should be constrained by experimental data in corrective loops to test precise hypotheses [14, p. 28]. Fig. 1.2 shows a simplified view of the (data-driven) scientific method life cycle. It distinguishes the phases of exploratory analytics (context of

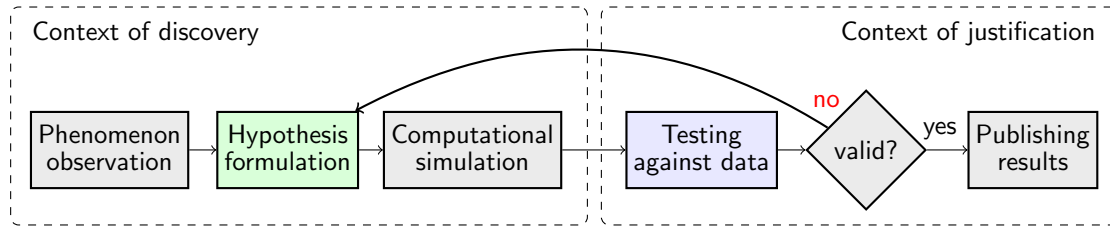


Figure 1.2. A view of the scientific method life cycle. It highlights hypothesis formulation and a backward transition to reformulation if predictions ‘disagree’ with observations.

discovery) and predictive analytics (context of justification), and highlights the loop between hypothesis formulation and testing [15].

Simulation data, being generated and tuned from a combination of theoretical and empirical principles, has a distinctive feature to be considered when compared to data generated by high-throughput technology in large-scale scientific experiments. It has a pronounced *uncertainty* component that motivates the use case of hypothesis data management for *predictive analytics* [10]. Essential aspects of hypothesis data management can be described in contrast to simulation data management as follows — Table 1.1 summarizes our comparison.

- *Sample data.* Hypothesis management shall not deal with the same volume of data as in simulation data management for exploratory analytics, but only samples of it. This is aligned, for example, with the architectural design of CERN’s particle-physics experiment and simulation ATLAS, where there are four tier/layers of data. The volume of data significantly decreases from (tier-0) the raw data to (tier-3) the data actually used for analyses such as hypothesis testing [8, p. 71-2]. Samples of raw simula-

Table 1.1. Simulation data management vs. hypothesis data management.

Simulation data management	Hypothesis data management
Exploratory analytics	Predictive analytics
Raw data	Sample data
Extremely large (TB, PB)	Very large (MB, GB)
Dimension-centered access pattern	Claim-centered access pattern
Denormalized for faster retrieval	Normalized for uncertainty factors
Batch-, incremental-only data updates	Probability distribution updates

tion data are to be selected for comparative studies involving competing hypotheses in the presence of evidence (sample observational data). This principle is also aligned with how data is delivered at model repositories. Since observations are usually less available, only the fragment (sample) of the simulation data that matches in coordinates the (sample) of observations is required out of simulation results for comparative analysis. For instance, we show in §6.2.2 a predictive analytical study extracted from the Virtual Physiological Rat Project (VPR1001-M) comparing sample simulation data (heart rates) from a baroreflex model with observations on a Dahl SS rat strain.⁵ The simulation is originally set to produce predictions in the time resolution of $t_{\Delta} = 0.01$. But since the observational sample is only as fine as $t_{\Delta} = 0.1$, there is no gain in rendering a predicted sample with $t_{\Delta} \geq 0.1$ for hypothesis testing. Note that such a ‘sampling’ does not incur in any additional uncertainty as typical of statistical sampling [16].

- *Claim-centered access pattern.* In simulation data management the access pattern is dimension-centered (e.g., based on selected space-time coordinates) and the data is denormalized for faster retrieval, as typical of Data Warehouses (DW’s) and OLAP applications.⁶ In particular, on account of the so-called ‘big table’ approach, each state of the modeled physical system is recorded in a large, single row of data. This is fairly reasonable for an Extract-Transform-Load (ETL) data ingesture pipeline characterized by batch-, incremental-only updates (see Fig. 1.3). Such a setting is in fact fit for exploratory analytics, as entire states of the simulated system shall be accessed at once (e.g., providing data to a visualization system). Altogether, data retrieval is critical and there is no risk of update anomalies. Hypothesis management, in contrast, should be centered on claims identified within the hypothesis structure w.r.t. available data dependencies. Since the focus is on resolving uncertainty for decision making (which

⁵ <http://virtualrat.org/computational-models/vpr1001/>.

⁶ On-Line Analytical Processing, as distinguished from OLTP (On-Line Transaction Processing). The latter is meant for transaction processing of daily queries and updates in operational systems, while the former is for analytical queries in Data Warehouses (DW’s) that gather a lot of data collected from different sources for decision making.

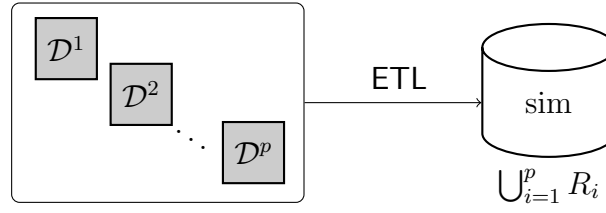


Figure 1.3. The usual data ingestion pipeline of simulation data management. Datasets $\bigcup_{i=1}^p \mathcal{D}^i$ generated by simulation trials on (hypothesis) models are loaded each into a ‘big’ table $\bigcup_{i=1}^p R_i$. The uncertainty is then “buried” in the database, which lacks a logical organization for enabling data-driven hypothesis management and predictive analytics.

hypothesis is a best fit?), the data must be normalized based on *uncertainty factors*. This is key for the correctness of uncertainty modeling and efficiency of probabilistic reasoning, say, in a probabilistic database [17, p.30-1].

- *Uncertainty modeling.* In uncertain and probabilistic data management [17], the uncertainty may come from two sources: *incompleteness* (missing data), and *multiplicity* (inconsistent data). Hypothesis management on sample simulation data is concerned with the multiplicity of prediction records due to competing hypotheses targeted at the same studied phenomenon. Such a multiplicity naturally gives rise to a probability distribution that may be initially uniform and eventually conditioned on observations. Conditioning is an applied *Bayesian inference* problem that translates into database update for transforming the prior probability distribution into a posterior [10].

Overall, hypothesis data management is also OLAP-like, yet markedly different from simulation data management.

A key point that distinguishes hypothesis management is that a fact or unit of data is defined by its **predictive content**. That is, every clear-cut predicted fact (w.r.t. available data dependencies) is a claim. Accordingly, the data should be decomposed and organized for a claim-centered access pattern.

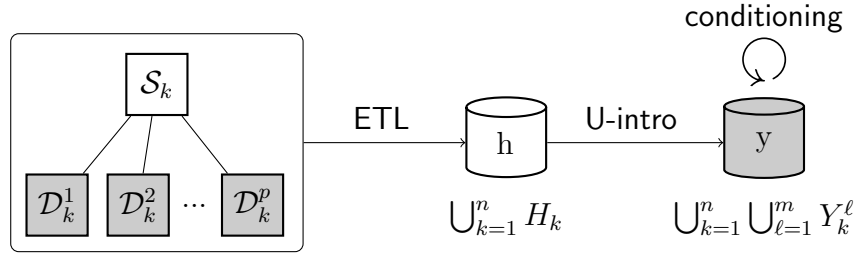


Figure 1.4. Pipeline for processing hypotheses as uncertain and probabilistic data. For each hypothesis k , its structure S_k is given in a machine-readable format, and all of its sample simulation data trials $\bigcup_{i=1}^p D_k^i$ are indicated their target phenomenon, say ϕ , to be loaded into a ‘big table’ H_k . Then the synthesis comes into play to read a base of possibly very many hypotheses $\bigcup_{k=1}^n H_k$ and transform them into a probabilistic database where each hypothesis is decomposed into claim tables $\bigcup_{\ell=1}^m Y_k^\ell$. A probability distribution is computed for each phenomenon ϕ , covering all the hypotheses and their trials targeted at ϕ . This distribution is then updated into a posterior in the presence of observational data.

To anticipate Chapter 2, the synthesis method we have developed in this thesis work for processing hypotheses as uncertain and probabilistic data comprises a design-theoretic pipeline (see Fig. 1.4) that extends the one shown in Fig. 1.3.

1.1.2 Probabilistic database design

Probabilistic databases (p-DB’s) have evolved into mature technology in the last decade with the emergence of new data models and query processing techniques [17]. One of the state-of-the-art probabilistic data models is the U-relational representation system with its probabilistic world-set algebra (p-WSA) implemented in MayBMS [18]. That is an elegant extension of the relational model we shall refer to in this thesis for the management of large-scale uncertain and probabilistic data.

We look at U-relations from the point of view of p-DB design, for which no formal design methodology has yet been proposed. Despite the advanced state of probabilistic data management techniques, a lack of methods for the systematic design of p-DBs may prevent wider adoption. The availability of design methods has been considered one of the key success factors for the rapid growth of applications in the field of Graphical Models (GM’s) [19], considered to inform research in p-DB’s [17, p. 14]. Analogously, we have proposed to distinguish methods for p-DB design in three classes [10]: (i) *subjective* construction, (ii) *learning* from

data, and (iii) *synthesis* from other kind of formal specification.

The first is the less systematic, as the user has to model for the data and correlations by steering all the p-DB construction process (MayBMS' use cases [18], e.g., are illustrated that way). The second comprises analytical techniques to extract the data and learn correlations from external sources, possibly unstructured, into a p-DB under some ad-hoc schema. This is the prevalent one up to date, motivated by information extraction and data integration applications [17, p. 10-3]. In this thesis we present a methodology of the third kind, as we extract data dependencies from some previously existing formal specification (the hypothesis mathematical structure) to synthesize a p-DB algorithmically. Such a type of construction method has been successful, e.g., for building Bayesian Networks [19]. To our knowledge, this thesis is the first synthesis method for p-DB design (cf. §5.6).

We shall develop means to extract the specification of a hypothesis and encode it into a U-relational DB for data-driven hypothesis management and analytics. That is, we shall flatten deterministic hypotheses into U-relations.

The synthesis method that we have developed for p-DB's relies on the extraction of *functional dependencies* (fd's; cf. [20, 21, 22]) that are basic input to algorithmic synthesis.⁷ For an example of fd, consider relation **FALL** in Fig. 1.1. There holds an fd $t \rightarrow v s$, meaning that values of attribute time t functionally determine values of both attributes velocity v and position s . More precisely, let μ and τ be any two tuples (rows) in an instance of relation (table) **FALL**. Then it satisfies fd $t \rightarrow v s$ iff $\mu[t] = \tau[t]$ implies $\mu[v s] = \tau[v s]$. In our illustrative relation **FALL**, that fd is, in particular, a *key constraint*, which means that (values of) t play the role of a key to (provide access to the values of) v and s in the relation.

A related concept which is also a major one for us is that of *normalization* [20, 21, 22], viz., to ensure that the DB resulting from a design process bears some desirable properties which are associated with some notion of *normal form* (ibid.). For hypothesis management, the uncertainty has to be modeled and should be normalized so that the uncertainty of one claim may not be undesirably mixed

⁷ In fact, it has been considered a critical failure in traditional DB design the lack of techniques to obtain important information such as fd's in the real world [23, p. 62].

with the uncertainty of another claim. It is expected to involve a processing of the causal dependencies implicit in the given hypothesis structure. We shall introduce in detail such concepts in context when necessary.

1.1.3 Structural equations

The flattening of the user mathematical models into hypothesis p-DB's, nonetheless, is not straightforward. It has been a goal of this thesis to investigate proper abstractions on mathematical models in order to (partly) capture their semantics, viz., to an extent that is tailored for hypothesis management (as opposed to, say, model solving). We shall abstract mathematical models into intermediary artifacts that are amenable to be further encoded into fd's.

In fact, given a system of equations with a set of variables appearing in them, in a seminal article Simon introduced an asymmetrical, functional relation among variables that establishes a (so-called) *causal ordering* [24]. That became known as *structural equation models* (SEM's) or just 'structural equations' (cf. also [25]). Along these lines, our goal is to extract the causal ordering implicit in the structure of a deterministic hypothesis into a set of fd's that guides our synthesis of U-relational DB's. As we shall see throughout this text,

the causal ordering we capture and process through fd's provides causal dependencies implicit in the predictive data that are very useful information to decompose uncertainty for the sake of probabilistic modeling and reasoning.

1.1.4 Uncertainty Model

In uncertain and probabilistic data management, there are essentially two sources of uncertainty: *incompleteness* (missing data), and *multiplicity* (inconsistent data).

The kind of uncertainty that is dealt with in this work is the multiplicity of hypothesis trial records identified to be targeted at the same phenomenon record. That is, the uncertainty arises from the existence of competing hypotheses. If multiple hypotheses and trials are inserted for the same phenomenon, the system interprets it as defining a probability distribution.

Such a probability distribution (usually uniform) on the multiplicity of competing hypotheses is in accordance with probability theory under possible-worlds semantics [17, Ch. 1]. It is modeled into the U-relational data model and its p-WSA operators, and implemented into the **MayBMS** system as we shall see in §5.1.⁸ The **conf()** aggregate operator, for instance, in spite of the name, performs standard (non-Bayesian) probabilistic inference on such probability distribution. Eventually, however, there is a need to condition the initial probability distribution in the presence of observations. For the conditioning, then, we shall adopt Bayesian inference so that the prior probability distribution can be updated to a posterior.

The informal discussion of this section opens the way for a number of technical research questions that we outline next.

- RQ5.** Is there an algorithm to, given a SEM, efficiently extract its causal ordering? What are the computational properties of this problem?
- RQ6.** What is the connection between SEM's and fd's? Can we devise an encoding scheme to 'orient equations' and then effectively transform one into the other with guarantees? Once we do it, what design-theoretic properties have such a set of fd's?
- RQ7.** Is such fd set ready to be used for p-DB schema synthesis as an encoding of the hypothesis causal structure? If not, what kind of further processing we have to do? Can we perform it efficiently by reasoning directly on the fd's? How does it relate to the SEM's causal ordering?
- RQ8.** Is the uncertainty decomposition required for predictive analytics reducible to the structure level (fd processing), or do we need to process the simulated data to identify additional uncertainty factors? Finally, what properties are desirable for a p-DB schema targeted at hypothesis management? Are they ensured by this synthesis method?
- RQ9.** Given all such a design-theoretic machinery to process hypotheses into (U-)relational DB's, what properties can we detect on the hypotheses back

⁸ Our own system of hypothesis management is to be delivered on top of the **MayBMS** backend.

at the conceptual level? Do we have now technical means to speak of hypotheses that are “good” in terms of principles of the philosophy of science?

The core of this thesis is devoted to answer these questions, and we shall accomplish it throughout Chapters 3, 4 and 5.

1.2. Thesis Statement

The statement of this thesis is that it is possible to effectively encode and manage large deterministic scientific hypotheses as uncertain and probabilistic data. Its key challenges are of both conceptual and technical nature. Conceptually, we provide core, non-obvious abstractions to define and encode hypotheses as data. Technically, we provide a number of algorithms that compose a design-theoretic pipeline to encode hypotheses as uncertain and probabilistic data, and verify their efficiency and correctness. The applicability and effectiveness of our method is demonstrated in realistic case studies in computational science.

Besides, it is worthwhile highlighting some non-goals of this thesis.

- N1.** Although we perform some sort of *information extraction* [26] for the acquisition of hypotheses from some model repositories on the web, it is very basic and ad-hoc in order to obtain a testbed for our method. That is, we are not proposing means for the systematic extraction of hypotheses from available sources. In fact we shall outline it in §7.3 as an important direction of future work.
- N2.** We do not address *solving* computational models or *numerical analytics* in any sense. In fact we rely on the numerical solvers (implemented into tools that we use) as ‘transaction processing’ systems, load their computed data into a relational ‘big’ fact table and then render it into U-relational tables synthesized by our method. We do not deal with *data visualization* either in any sense.
- N3.** The efficiency and scalability of *query processing* in p-DB’s, in particular U-relational’s MayBMS and its p-WSA (which we rely on) is *not* addressed or

evaluated in this thesis. In fact, the performance of U-relations and p-WSA has been extensively evaluated and shown to be effective [27, 18]. All performance tests carried out in this thesis comprise our design-theoretic techniques for the encoding and synthesis of U-relational hypothesis databases.

- N4.** In terms of *uncertainty* and *statistical analysis*, we stick to (i) process some well-defined forms of multiplicity in the data which constitute the model of uncertainty dealt with in this work; then (ii) by relying on **MayBMS** we perform probabilistic inference; and (iii) eventually (at application level) we perform Bayesian inference and so that a posterior probability distribution is propagated through p-DB updates. We do not provide any additional form of *uncertainty management*. Rather, we manage the data extracted into the system (under user control) and process its uncertainty in terms of the specific sources of uncertainty recognized in Υ -DB (cf. Chapter 2).

1.3. Thesis Contributions

The contributions of this thesis are outlined as follows.

1.3.1 Innovative Contributions

This thesis presents the vision of hypotheses as data (and its use case) so-called Υ -DB vision. It has been published in the vision track of VLDB 2014 [10], for its (sic.) potentially high-impact visionary content. The innovative system of Υ -DB has been described in a ‘system prototype demonstration’ paper [28].⁹

1.3.2 Technical Contributions

This thesis presents specific technical developments over the Υ -DB vision. In short, it shows how to encode deterministic hypotheses as uncertain and probabilistic data. Our detailed technical contributions (cf. Chapters 3, 4, and 5) are formulated into a formal method for the design of hypothesis p-DB’s which is described in a technical report [29].¹⁰ The method, together with our realistic testbed scenarios and performance evaluation, are yet to be published.

⁹ Preliminary version available at CoRR abs/1411.7419.

¹⁰ Preliminary version available at CoRR abs/1411.5196.

1.4. Thesis Outline

The structure of the remainder of this thesis is outlined for reference.

Chapter 2. [Υ -DB Vision]. The research vision of hypotheses as (uncertain and probabilistic) data, the characterization of its use case, key points and technical challenges are presented.

Chapter 3. [Encoding]. The problem of encoding a hypothesis ‘as data’ given its formal specification (set of mathematical equations) is presented and addressed by an encoding scheme that transforms the equations into fd’s with guarantees in terms of preserving the hypothesis causal structure.

Chapter 4. [Causal Reasoning]. It is presented a technique for causal reasoning as acyclic pseudo-transitive reasoning over the encoded fd’s. It processes the hypothesis causal ordering to find the ‘first causes’ for each of its predictive variables.

Chapter 5. [p-DB Synthesis]. It is presented a technique to address the problem of uncertainty introduction and propagation for the transformation of hypotheses into U-relational databases. The synthesized U-database is shown to bear desirable properties for hypothesis management and predictive analytics.

Chapter 6. [Applicability]. A discussion of applicability, the implementation of the proposed techniques into a prototype system for test and demonstration of the vision realization through realistic case studies are presented.

Chapter 7. [Conclusions]. Research questions are revisited, and the significance and limitations of the thesis with directions to future work and final considerations are discussed.

Chapter 2

Vision: Hypotheses as Data

High-throughput technology and large-scale scientific experiments provide scientists with *empirical* data that has to be extracted, transformed and loaded before it is ready for analysis [1]. In this vision we consider *theoretical* data, or data generated by simulation from deterministic scientific hypotheses, which also needs to be pre-processed to be analyzed.

Hypotheses as data. In view of the age of data-driven science, we consider deterministic scientific hypotheses from a multi-fold point of view: formed as principles or learned in large scale,¹ hypotheses are formulated mathematically and coded in a program that is run to give their *decisive* form of data (see Fig. 1.1).

Uncertain data. The semantic structure of relation FALL (Fig. 1.1), item (iv) can be expressed by the functional dependency (fd) $t \rightarrow v s$. This is typical semantics assigned to empirical data in the design of experiment databases. A space-time dimension (like time t in our example) is used as a key to observables (like velocity v and position s). In *empirical* uncertainty, it is such “physical” dimension keys like t that may be violated, say, by alternative sensor readings.

Hypotheses, however, are tentative explanations of phenomena [15], which characterizes a different kind of uncertain data. In order to manage such *theoretical* uncertainty, we shall need two special attributes to compose, say, the epistemological dimension of keys to observables: ϕ , identifying the studied phenomena; and v , identifying the hypotheses aimed at explaining them. That is, we shall leverage the semantics of relations like FALL to $\phi v t \rightarrow v s$. This leap is a core abstraction

¹ As exhibited, e.g., in the Eureka project [2].

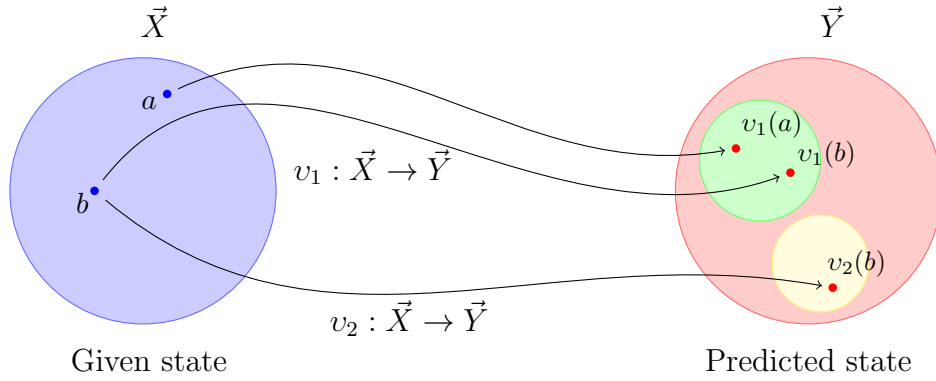


Figure 2.1. Deterministic scientific hypotheses seen as alternative functions to predict data, giving rise to both theoretical and empirical sources of uncertainty.

in this vision of Υ -DB.

Predictive data. Scientific hypotheses are tested by way of their predictions [15]. In the form of mathematical equations, hypotheses symmetrically relate aspects of the studied phenomenon. However, for computing predictions, deterministic hypotheses are applied asymmetrically as *functions* [30]. They take a given valuation over input variables (parameters) to produce values of output variables (predictions). By observing that, we shall seek a principled method to transform the (symmetric) mathematical equations of a hypothesis into (asymmetric) fd's.

By looking at deterministic hypotheses as alternative functions to predict data (see Fig. 2.1), in this vision we shall deal with two sources of uncertainty. Given a well-defined context with a set of alternative hypotheses aimed at explaining (providing predictions for) a selected phenomenon:

- *Theoretical* uncertainty,² comprises selecting the best tentative model (function) to produce (the best) data?
- *Empirical* uncertainty,³ comprises, for each candidate model, what is the (parameter) input setting that calibrates it the best way for the selected phenomenon?

Note that these two sources of uncertainty are intertwined in that one cannot 'clean' one without cleaning the other — neither theory nor parameters are directly

² That is, multiplicity of hypothesis entries associated with a phenomenon.

³ That is, multiplicity of hypothesis trial entries associated with a phenomenon.

observable, but only their joint results (the predictions) [15]. In this thesis we aim at providing means to support such kind of ‘integrated’ analytics.

Applications. Big computational science research programs such as the Human Brain Project,⁴ or Cardiovascular Mathematics,⁵ are highly-demanding applications challenged by such theoretical (big) data. Users need to analyze results of hundreds to thousands of data-intensive simulation trials.

Besides, recent initiatives on web-based model repositories have been fostering large-scale model integration, sharing and reproducibility in the computational sciences (e.g., [31, 32, 33]). They are growing reasonably fast on the web, (i) promoting some MathML-based standard for model specification, but (ii) with limited integrity and lack of support for rating/ranking competing models. For those two reasons, they provide a strong use case for our vision of hypothesis management. The Physiome project [33, 34], e.g., is planned to integrate several large deterministic models of human physiology — a fairly simple model of the human cardiovascular system, e.g., has about 600+ variables.

Also, there is a pressing call for *deep* predictive analytic tools to support users assessing what-if scenarios in business enterprises [35]. Deep predictive analytics are based on first principles (deterministic hypotheses) and go beyond descriptive analytics or shallow predictive analytics such as statistical forecasting (ibid.).⁶

U-relations. All that ratifies that *hypothesis management* is a promising class of applications for probabilistic DB’s. The vision of Υ -DB is currently set to be delivered on top of U-relations and probabilistic world-set algebra (p-WSA) [18]. These were developed in the influential MayBMS project.⁷ As implied by some of its design principles, viz., compositionality and the ability to introduce uncertainty, MayBMS’ query language fits well to hypothesis management. We shall look at it, as previously mentioned, from the point of view of a synthesis method for p-DB design. We shall particularly make use of the **repair key** operation, which gives

⁴ <http://www.humanbrainproject.eu/>.

⁵ <http://icerm.brown.edu/tw14-1-pdecn>.

⁶ The concept of ‘deep’ predictive analytics is from Haas et al. [35], and is discussed in more detail in §2.6.1.

⁷ Project website: <http://maybms.sourceforge.net/>. MayBMS is as a backend extension of PostgreSQL. It offers all the traditional querying capabilities of the latter in addition to the uncertain and probabilistic’s.

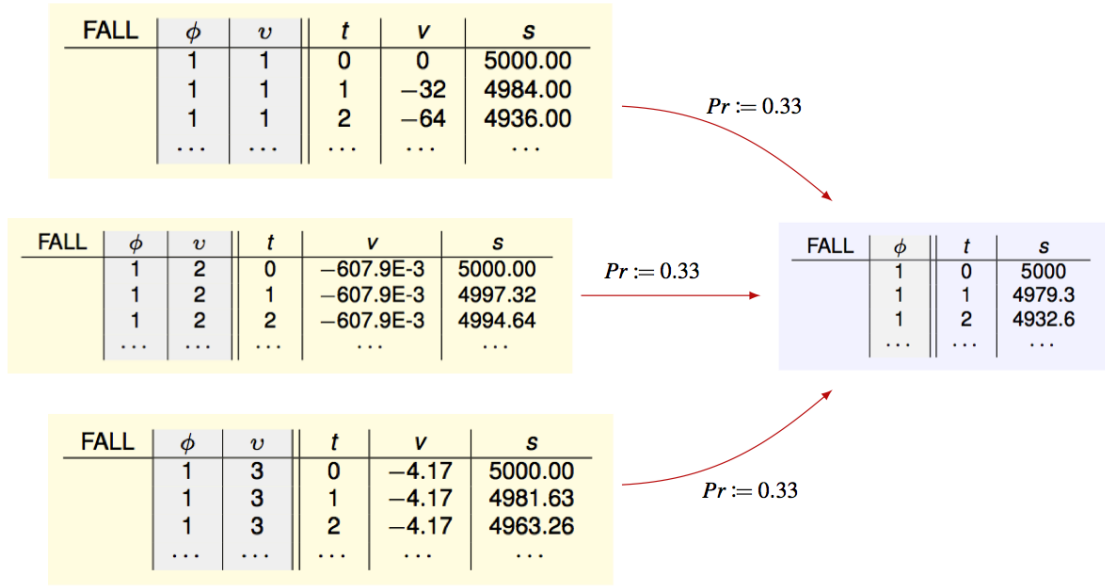


Figure 2.2. Predictive analytics in a data-intensive hypothesis evaluation study: hypotheses (simulated data) compete to explain a phenomenon (observed data).

rise to alternative worlds as maximal-subset repairs of an argument key.

Predictive analytics tool. In database research (e.g., [36]), uncertainty is usually seen as an undesirable property that hinders data quality. We shall refer to U-relations and p-WSA as implemented in **MayBMS**, nonetheless, to show that the ability to introduce *controlled* uncertainty into an (otherwise complete) simulation dataset can be a tool for ‘deep’ predictive analytics on a set of competing or alternative hypotheses. Fig.2.2 shows such a scenario of hypotheses ‘as data’ compete to explain a phenomenon ‘as data.’

As a roadmap to most of the remainder of this chapter, we claim that if hypotheses can be encoded and identified (see §2.2), and their uncertainty quantified by some probability distribution (see §2.4), then they can be rated/ranked and browsed by the user under selectivity criteria. Furthermore, their probabilities can be conditioned for possibly being re-ranked in the presence of evidence (see §2.5).

2.1. Running Example

Let us consider Example 1 for the presentation of the vision.

Example 1 A research is conducted on the effects of gravity on a falling object in the Earth’s atmosphere. Scientists are uncertain about the precise object’s

density and its predominant state as a fluid or a solid. Three hypotheses are then considered as alternative explanations of the fall (see Fig. 2.3). Due to parameter uncertainty, six simulation trials are run for \mathcal{H}_1 , and four for \mathcal{H}_2 and \mathcal{H}_3 each. \square

PHENOMENON	ϕ	Description
	1	Effects of gravity on an object falling in the Earth's atmosphere.

HYPOTHESIS	v	Name
	1	Law of free fall
	2	Stokes' law
	3	Velocity-squared law

Figure 2.3. Descriptive (textual) data of Example 1.

H_1	tid	ϕ	v	t	g	v_0	s_0	a	v	s
	1	1	1	0	32.0	0	5000	-32.0	0	5000
	2	1	1	0	32.0	10	5000	-32.0	10	5000
	3	1	1	0	32.0	20	5000	-32.0	20	5000
	4	1	1	0	32.2	0	5000	-32.2	0	5000
	5	1	1	0	32.2	10	5000	-32.2	10	5000
	6	1	1	0	32.2	20	5000	-32.2	20	5000
	1	1	1	0.1	32.0	0	5000	-32.0	-3.2	4999.84
	2	1	1	...	32.0	0	5000	-32.0

Figure 2.4. ‘Big’ fact table H_1 of hypothesis $v = 1$ loaded with simulation raw data: trials on H_1 are identified by **tid**.

The construction of Υ -DB, a Data Warehouse (DW), requires a simple user description of a research. That is, descriptive records of the phenomena and hypotheses dimensions (see Fig. 2.3) are to be inserted first such that basic referential constraints are satisfied by their associated datasets (fact tables). For instance, each one of the six trial datasets for hypothesis \mathcal{H}_1 shall reference its id $v = 1$ as a foreign key from table HYPOTHESIS further in their synthesized relations.

Fig. 2.4 shows the ‘big’ fact table H_1 for hypothesis $v = 1$ loaded with its trial datasets for phenomenon $\phi = 1$. Although table H_1 is denormalized for faster data retrieval as usual in DW’s, the extraction of the hypothesis equations allows to render it automatically since all variables must appear in some equation. Now we proceed to the hypothesis encoding and start to address research questions RQ1-4.

2.2. Hypothesis Encoding

We aim at extracting, for each hypothesis, a set of fd's from its mathematical equations. Suppose we are given a set of equations of hypothesis \mathcal{H}_1 below, and let us examine the set Σ_1 of fd's we target at.⁸

\mathcal{H}_1 . Law of free fall	
$a(t)$	$= g$
$v(t)$	$= -gt + v_0$
$s(t)$	$= -(g/2)t^2 + v_0 t + s_0$

$$\Sigma_1 = \{ \begin{array}{l} \phi \rightarrow g, \\ \phi \rightarrow v_0, \\ \phi \rightarrow s_0, \\ g v \rightarrow a, \\ g v_0 t v \rightarrow v, \\ g v_0 s_0 t v \rightarrow s \end{array} \}.$$

In order to derive Σ_1 from the equations of \mathcal{H}_1 , we focus on their implicit data dependencies and get rid of constants and possibly complex mathematical constructs. Equation $v(t) = -gt + v_0$, e.g., written this way (roughly speaking), suggests that v is a prediction variable *functionally dependent* on t (the physical dimension), g and v_0 (the parameters). Yet a dependency like $g v_0 t \rightarrow v$ may hold for infinitely many equations.⁹ In fact, we need a way to identify \mathcal{H}_1 's mathematical formulation precisely, i.e., an abstraction of its data-level semantics. This is achieved by introducing hypothesis id v as a special attribute in the fd (see Σ_1).

This is a *data representation* of a deterministic scientific hypothesis. It is built into an encoding scheme (see §3.4) that leverages the semantics of structural equations.

The other special attribute, the phenomenon id ϕ , is supposed to be a key to the values of parameters, i.e., determination of parameters is an empirical, phenomenon-dependent task. The fd $\phi \rightarrow g v_0 s_0$ is to be (expectedly) violated when the user

⁸ Recall that a rigorous presentation of the method to encode fd set Σ_1 is due by Chapter 3.

⁹ Think of, say, how many polynomials satisfy that dependency 'signature.'

is uncertain about the values of parameters. The same rationale applies to derive $\Sigma_2 = \Sigma_3$ from the equations of \mathcal{H}_2 , \mathcal{H}_3 below. These, n.b., vary in structure w.r.t. \mathcal{H}_1 (e.g., they include parameter D , the object's diameter).

\mathcal{H}_2 . Stokes' law	\mathcal{H}_3 . Velocity-squared law
$a(t) = 0$	$a(t) = 0$
$v(t) = -\sqrt{gD/4.6 \times 10^{-4}}$	$v(t) = -gD^2/3.29 \times 10^{-6}$
$s(t) = -t\sqrt{gD/4.6 \times 10^{-4}} + s_0$	$s(t) = -(gD^2/3.29 \times 10^{-6})t + s_0$

$$\Sigma_2 = \Sigma_3 = \left\{ \begin{array}{l} \phi \rightarrow g, \\ \phi \rightarrow D, \\ \phi \rightarrow s_0, \\ \phi \rightarrow a, \\ gDv \rightarrow v, \\ gDs_0tv \rightarrow s \end{array} \right\}.$$

The key point here is that, if the hypothesis structure (set of equations) is given in a machine-readable format for mathematics, then the method to extract the hypothesis fd set from its equations can be carefully designed based on such hypothesis data representation abstraction. In fact, we shall explore W3C's MathML as a format for hypothesis specification.¹⁰

2.3. Reasoning over FD's

Once each hypothesis fd set has been extracted, some reasoning is to be performed to discover implicit data dependencies. In fact, dependency theory is equipped with a formal system (cf. §4.1) for reasoning over fd sets like Σ_1 and derive other fd's in its closure Σ_1^+ . As we elaborate on in Chapter 4, we shall be particularly concerned with the pseudo-transitivity inference rule. Applied over fd's $\{\phi \rightarrow g, g v \rightarrow a\} \subset \Sigma_1$, for instance, it gives us $\phi v \rightarrow a$. This inference allows us to observe that $\{g\}$ is a 'factor' on the uncertainty of a , but $\{\phi v\}$ should be a dimensional *key constraint* for values of a .

In fact, note that derived fd's like $\langle \phi v, a \rangle \in \Sigma_1^+$, which should be a constraint on values of a in H_1 , are (expectedly) violated in the presence of uncertainty: ob-

¹⁰ <http://www.w3.org/Math/>.

serve in Fig. 2.4 the multiplicity $\{32.0, 32.2\}$ of values of a under the same pair $(\phi \mapsto 1, v \mapsto 1)$, which should functionally determine them in H_1 . For that reason we admit a special attribute ‘trial id’ tid to be overimposed into H_1 for a trivial repair, provisionally, until uncertainty can be introduced in a controlled way by synthesis ‘4U.’ It is meant to identify simulation trials and “pretend” certainty not to lose the integrity of the data. It is under this imposed certainty that the raw simulation trial data is safely loaded from files (see Fig. 2.4). Note, however, how ‘certainty’ is held at the expense of redundancy and, mostly important, opaqueness for predictive analytics (since tid isolates or hides the inconsistency w.r.t. to the violated constraints). This is until the next stage of the Υ -DB construction pipeline, when uncertainty is to be introduced in a controlled manner.

2.4. Uncertainty Introduction

Before we proceed to the uncertainty introduction procedure, note in relation H_1 (Fig. 2.4), that the predicted acceleration values a are such that an association between the hypothesis and a target phenomenon, viz., $(\phi \mapsto 1, v \mapsto 1)$ is established. In fact, as of the insertion of each hypothesis trial dataset, the user must set for it a target phenomenon. This may be non-obvious but is quite convenient a design decision for the envisioned system of Υ -DB because hypotheses, as (abstract) universal statements [15], can only be derived predictions from (be empirically grounded) by assigning (calibrating) them onto some real-world phenomenon. This assignment is set at data entry time because in fact it only holds at the data level.¹¹ It is to be recorded in an ‘explanation’ table named H_0 by default (see Fig. 2.5, top), being provided with weights for establishing a prior probability distribution which (by user choice) may or may not be uniform.

The data transformation of ‘certain’ to ‘uncertain’ relations then starts with query Q_0 , whose result set is materialized into U-relational table Y_0 (see Fig. 2.5). As we introduce in detail in §5.1, U-relations have in their schema a set of pairs (V_i, D_i) of *condition columns* [18] to map each discrete random variable x_i created by the **repair-key** operation to one of its possible values (e.g., $x_0 \mapsto 1$). The

¹¹ Hypotheses are ‘universal’ by definition [15]. They (must) qualify for a *class* of different situated phenomena, while its predictive datasets must be very specific (for one specific situation).

H_0	ϕ	v	Conf	Y_0	$V \mapsto D$	ϕ	v	W	$V \mapsto D$	Pr
	1	1	3		$x_0 \mapsto 1$	1	1		$x_0 \mapsto 1$.6
	1	2	1		$x_0 \mapsto 2$	1	2		$x_0 \mapsto 2$.2
	1	3	1		$x_0 \mapsto 3$	1	3		$x_0 \mapsto 3$.2

Figure 2.5. ‘Explanation’ relational table H_0 and its associated U-relational table Y_0 rendered by application of the **repair-key** operation.

world table W is internal to **MayBMS**’ and automatically stores their marginal probabilities. The formal semantics the **repair-key** operation is given in §5.1.

Q_0 . **create table** Y_0 **as select** ϕ, v **from** (**repair key** ϕ **in** H_0 **weight by** Conf);

The possible-world semantics of p-DB’s (cf. §5.1) can be seen as a generalization of *data cleaning*. In the context of p-DB’s [17], data cleaning does not have to be one-shot — which is more error-prone [37]. Rather, it can be carried out gradually, viz., by keeping all mutually inconsistent tuples under a probability distribution (ibid.) that can be updated in face of evidence until the probabilities of some tuples eventually tend to zero to be eliminated. This motivates Remark 1.

Remark 1 Consider U-relational table Y_0 (Fig. 2.5). Note that it abstracts the goal of a data-intensive hypothesis evaluation study, or the scientific method itself [15], as the repair of each ϕ as a key. That is, in Υ -DB users can develop their research directly upon data with support of query and update capabilities to rate/rank their hypotheses v w.r.t. each ϕ , until the relationship $r(\phi, v)$ is repaired to be a *function* $f: \Phi \rightarrow \Upsilon$ from each phenomenon ϕ to its best explanation v . \square

Given a ‘big’ fact table such as H_1 , we need to identify/group the correlated input attributes under independent uncertainty units, viz., ‘u-factors,’ each one associated with a random variable.¹² We illustrate that by means of query Q_1 , which materializes view $Y_1[g]$ for (let $g = Z_i$) identified u-factor $Z_i \subseteq Z$ in $H_1[\phi, Z]$.

¹² An attribute can be inferred ‘input’ (viz., a parameter) by means of fd reasoning (cf. §3.4).

$H_1[\text{tid}, \phi, Z]$	tid	ϕ	g	v_0	s_0
	1	1	32	0	5000
	2	1	32	10	5000
	3	1	32	20	5000
	4	1	32.2	0	5000
	5	1	32.2	10	5000
	6	1	32.2	20	5000

$Y_1[g]$	$V \mapsto D$	ϕ	g
	$x_1 \mapsto 1$	1	32
	$x_1 \mapsto 2$	1	32.2

W	$V \mapsto D$	Pr
	\dots	\dots
	$x_1 \mapsto 1$.5
	$x_1 \mapsto 2$.5

Figure 2.6. Result set of query Q_1 on simulation trial dataset for hypothesis H_1 .

Q_1 . **create table** $Y_1[g]$ **as select** $U.\phi, U.g$ **from** (**repair key** ϕ in
 (**select** $\phi, g, \text{count}(\ast)$ **as** Fr **from** H_1 **group by** ϕ, g)
weight by Fr) **as** U;

The result set of Q_1 is stored in $Y_1[g]$, see Fig. 2.6. Note that the possible values of g are mapped to random variable x_1 , and that table H_1 is considered source for a joint probability distribution (on the values of H_1 's input parameters) which may not be uniform: we count the frequency Fr of each possible value of a u-factor $Z_i \subseteq Z$ (as done for g in Q_1) and pass it as argument to the **weight-by** construct.

So far, we have presented informally the procedure of *u-factorization*. Now we proceed to *u-propagation* — both are presented rigorously in Chapter 5. We consider $g v \rightarrow a \in \Sigma_1$ again in order to synthesize predictive U-relation $Y_1[a]$. Since a is functionally determined by v and g only, and these are independent, we propagate their uncertainty onto a into $Y_1[a]$ by query Q_2 .

Q_2 . **create table** $Y_1[a]$ **as select** $H_1.\phi, H_1.v, H_1.a$ **from** $H_1, Y_0, Y_1[g]$ **as** G
where $H_1.\phi=Y_0.\phi$ **and** $H_1.v=Y_0.v$ **and** $G.\phi=H_1.\phi$ **and** $G.g=H_1.g$;

Query Q'_2 (not shown) then selects ϕ, v and a from $Y_i[a]$ for each $i = 1..3$. The result sets of Q_2 and Q'_2 (resp. $Y_1[a]$ and $Y[a]$) are shown in Fig. 2.7.

$Y_1[a]$	$V_0 \mapsto D_0$	$V_1 \mapsto D_1$	ϕ	v	a
	$x_0 \mapsto 1$	$x_1 \mapsto 1$	1	1	-32
	$x_0 \mapsto 1$	$x_1 \mapsto 2$	1	1	-32.2

$Y[a]$	$V_0 \mapsto D_0$	$V_1 \mapsto D_1$	ϕ	v	a
	$x_0 \mapsto 1$	$x_1 \mapsto 1$	1	1	-32
	$x_0 \mapsto 1$	$x_1 \mapsto 2$	1	1	-32.2
	$x_0 \mapsto 2$	—	1	2	0
	$x_0 \mapsto 3$	—	1	3	0

Figure 2.7. U-relational predictive tables rendered by query using the fd's.

Compare relations $H_1[a]$ and $Y_1[a]$. By accounting for the correlations captured in the fd $g v \rightarrow a$, we could propagate onto a the uncertainty coming from the hypothesis and the only parameter a is sensible to, thus precisely situating tuples of $Y_1[a]$ in the space of possible worlds. The same is done for predictive attributes v and s . In the end, Υ -DB shall be ready for predictive analytics, i.e., with all competing predictions as possible alternatives which are mutually inconsistent.

A key point here is that all the synthesis process is amenable to algorithm design. Except for the user ‘research’ description, the Υ -DB construction is fully automated based on the hypothesis structure (set of equations) and the raw hypothesis trial data.

2.5. Predictive Analytics

Users of Example 1, has to be able, say, to query phenomenon $\phi = 1$ w.r.t. predicted position s at specific values of time t by considering all hypotheses v admitted. That is illustrated by query Q_3 , which creates integrative table $Y[s]$; and by query Q_4 , which computes the *confidence* aggregate operation [18] for all s tuples where $t = 3$ (Fig. 2.8 shows Q_4 's result, apart from column **Posterior**).

The confidence on each hypothesis for the specific prediction of Q_4 is split due to parameter uncertainty such that they sum up back to its total confidence. For H_2 and H_3 , e.g., we have $\{g D s_0 t v \rightarrow s\} \subset \Gamma$, where $\Gamma = \Sigma_2 = \Sigma_3$. Since g and D are the parameter uncertainty factors of s (s_0 is certain), with 2 possible values (not shown) each, then there are only $2 \times 2 = 4$ possible s tuples for H_2 and H_3 each. Considering all hypotheses v for the same phenomenon ϕ , the confidence

values sum up to one in accordance with the laws of probability.

```
Q3. create table Y[s] as select U.ϕ, U.v, U.t, U.s from
      (select ϕ, v, t, s from Y1[s] union all
      select ϕ, v, t, s from Y2[s] union all
      select ϕ, v, t, s from Y3[s]) as U, Y0
      where U.ϕ=Y0.ϕ and U.v=Y0.v;
```

```
Q4. select ϕ, v, s, conf() as Prior from Y[s] where t=3
      group by ϕ, v, s order by Prior desc;
```

Y[s]	ϕ	v	s	Prior	Posterior
	1	1	2188.36	.1	.167
	1	1	2205.82	.1	.168
	1	1	2320.51	.1	.167
	1	1	2337.97	.1	.165
	1	1	2452.66	.1	.149
	1	1	2470.12	.1	.145
	1	2	2930.59	.05	.020
	1	2	2943.44	.05	.019
	1	2	4991.92	.05	.000
	1	2	4991.97	.05	.000
	1	3	4778.87	.05	.000
	1	3	4779.56	.05	.000
	1	3	4944.72	.05	.000
	1	3	4944.89	.05	.000

Figure 2.8. Analytics on predicted position s conditioned on observation.

Users can make informed decisions in light of such confidence aggregates, which are to be eventually conditioned in face of evidence (observed data). Example 2 features such kind of Bayesian conditioning for *discrete* random variables mapped to the possible values of predictive attributes (like position s) whose domain are *continuous*.

Example 2 Suppose position $s = 2250$ feet is observed at $t = 3$ secs, with standard deviation $\sigma = 20$. Then, by applying Bayes' theorem for normal mean with a discrete prior [16], Prior is updated to Posterior (see Fig. 2.8). \square

The procedure uses normal density function (2.1), with (say) $\sigma = 20$, to get the likelihood $f(y | \mu_k)$ of each alternative prediction of s from $Y[s]$ as mean μ_k given y at observed $s = 2250$. Then it applies Bayes' rule (2.2) to get the posterior $p(\mu_k | y)$ [16].

$$f(y | \mu_k) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(y-\mu_k)^2} \quad (2.1)$$

$$p(\mu_k | y) = f(y | \mu_k) p(\mu_k) / \sum_{i=1}^n f(y | \mu_i) p(\mu_i) \quad (2.2)$$

In the general case (cf. examples shown in Chapter 6), we actually have phenomenon ‘as data:’ a sample of independent observed values y_1, \dots, y_n (e.g., Brazil’s population observed by census over the years). Then, the likelihood $f(y_1, \dots, y_n | \mu_k)$ for each competing trial μ_k , is computed as a product $\prod_{j=1}^n f(y_j | \mu_{kj})$ of the single likelihoods $f(y_j | \mu_{kj})$ [16]. Bayes’ rule is then settled by (2.3) to compute the posterior $p(\mu_k | y_1, \dots, y_n)$ given prior $p(\mu_k)$.

$$p(\mu_k | y_1, \dots, y_n) = \frac{\prod_{j=1}^n f(y_j | \mu_{kj}) p(\mu_k)}{\sum_{i=1}^m \prod_{j=1}^n f(y_j | \mu_{ij}) p(\mu_i)} \quad (2.3)$$

As a result, the prior probability distribution assigned to u-factors via **repair key** is to be eventually conditioned on observed data. This is an applied *Bayesian inference* problem that translates into a *p-DB update* one to induce effects of posteriors back to table W . In a first prototype of the Υ -DB system (cf. 6.3), we accomplish it by performing Bayesian inference at application level and then applying p-WSA’s **update** (a variant of SQL’s **update**) into **MayBMS**. This solution is good enough to let us complete use case demonstrations of Υ -DB.¹³

2.6. Related Work

The vision of managing hypotheses as data has some roots in Porto and Spaccapietra [38], who motivated a conceptual data model to support (the so-called) *in silico* science by means of a scientific model management system. We

¹³ Cf. Chapter 6, and §6.3 in particular.

discuss now the work we understand to be mostly related to our vision of data-driven hypothesis management and analytics.

2.6.1 Models-and-data

Haas et al. [35] provide an original long-term perspective on the evolution of database technology. They characterize the data typically managed by traditional DB systems as a record about the past, not a conclusion or an insight or a solution (ibid.). In the context of scientific databases, e.g., their position is suggestive that DB technology has been designed for empirical data, not the theoretical data generated by simulation from domain-specific principles or scientific hypotheses.

They recognize current DB technology to have raised the art of scalable ‘descriptive’ analytics to a very high level; but point out, however, that nowadays (sic.) what enterprises really need is ‘prescriptive’ analytics to identify optimal business, policy, investment, and engineering decisions in the face of uncertainty. Such analytics, in turn, shall rest on *deep* ‘predictive’ analytics that go beyond mere statistical forecasting and are imbued with an understanding of the fundamental mechanisms that govern a system’s behavior, allowing what-if analyses [35]. In sum, there is a pressing call for *deep* predictive analytic tools in business enterprises as much as in science’s.

In comparison with the Υ -DB vision, Haas et al. are proposing a long-term *models-and-data* research program to pursue data management technology for deep predictive analytics. They discuss strategies to extend query engines for model execution within a (p-)DB. Along these lines, query optimization is understood as a more general problem with connections to algebraic solvers.

Our framework in turn essentially comprises an abstraction and technique for the encoding of *hypotheses as data*. It can be understood (in comparison) as putting models strictly into a flattened data perspective. For that reason it has been directly applicable by building upon recent work on p-DBs [17]. In principle, it can be integrated into, say, the OLAP layer of the models-and-data project.

2.6.2 Scientific simulation data

As previously mentioned, science’s ETL is distinguished by its unfrequent, incremental-only updates and by having large raw files as data sources [8]. Challenges for

enabling an efficient access to high-resolution, raw simulation data have been documented from both supercomputing,[6] and database research viewpoints;[39] and pointed as key to the use case of *exploratory analytics*. The extreme scale of the raw data has motivated such non-conventional approaches for data exploration, viz., the ‘immersive’ query processing (move the program to the data) [6, 40], or ‘in situ’ query processing in the raw files [41, 42]. Both exploit the spatial structure of the data in their indexing schemes.

That line of research is motivated for equipping scientist end-users for an immediate interaction with their very large simulation datasets.¹⁴ The NoDB approach, in particular, argues to eliminate such ETL phase (viz., the loading) for a direct access to data ‘in situ’ in the raw data files [42]. In fact, data exploration is a fundamental use case of data-driven science.

Nonetheless, being generated from first principles or learned deterministic hypotheses, simulation data has a pronounced uncertainty component that motivates a another use case, viz., the case of hypothesis management and *predictive analytics* [35, 10]. As we have motivated in §1.1, the latter requires probabilistic DB design for enabling uncertainty decomposition (factorization).

Hypothesis management shall not deal with the same volume of data as in simulation data management for exploratory analytics, but samples of it (cf. Table 1.1 for a comparison). For instance, in CERN’s particle-physics experiment ATLAS there are four tier/layers of data management. The volume of data significantly decreases from the (tier-0) raw data to the (tier-3) data actually used for analyses such as hypothesis testing [8, p. 71-2].

Overall, the overhead incurred in loading samples of raw simulation trial datasets into a p-DB is justified for enabling a principled hypothesis evaluation and rating/ranking according to the scientific method.

2.6.3 Hypothesis encoding

Our framework is comparable with Bioinformatics’ initiatives that address hypothesis encoding into the RDF data model [43]: (i) the Robot Scientist [44] is a knowledge-base system (KBS) for automated generation and testing of hypotheses

¹⁴ Sometimes phrased ‘here is my files, here is my queries, where are my results?’ [41].

about what genes encode enzymes in the yeast organism; (ii) HyBrow [45] is a KBS for scientists to test their hypotheses about events of the galactose metabolism also of the yeast organism; and (iii) SWAN [46] is a KBS for scientists to share hypotheses on possible causes of the Alzheimer disease.

The Robot Scientist relies on rule-based logic programming analytics to automatically generate and test RDF-encoded hypotheses of the kind ‘gene G has function A’ against RDF-encoded empirical data [44]. HyBrow is likewise, but hypotheses are formulated by the user about biological events [45]. SWAN in turn disfavours analytic techniques for hypothesis evaluation and focus on descriptive aspects: hypotheses are high-level natural language statements retrieved from publications. Each ‘hypothesis’ is associated with lower-level ‘claims’ (both RDF-encoded) that are meant to support it on the basis of some empirical evidence (RDF-encoded gene/protein data). In particular, SWAN [46] differs from the former in that each hypothesis is unstructured, being then more related to efforts on the retrieval of textual claims from the narrative fabric of scientific reports [47].

All of them though, consist in some ad-hoc RDF encoding of sequence and genome analysis hypotheses under varying levels of structure (viz., from ‘gene G has function A’ statements to free text). Our framework in turn consists in the U-relational encoding of hypotheses from mathematical equations, which is (to our knowledge) the first work on hypothesis relational encoding.

Finally, as for hypothesis evaluation and comparison analytics, the Υ -DB vision is distinguished in terms of its Bayesian inference approach. The latter has been pointed out as a major direction for the improvement of the Bioinformatics’ initiatives just mentioned (cf. [43, p. 13]), and is in fact an influential model of decision making for hypothesis evaluation [15, p. 220].

2.7. Summary: Key Points

We outline some key points in the Υ -DB vision:

- ‘Structured deterministic hypotheses’ are encoded as theoretical data and distinguished from empirical data by the introduction of an *epistemological dimension* into their semantic structure.

- Two *sources of uncertainty* are considered: theoretical uncertainty, originating from competing hypotheses; and empirical uncertainty, derived from alternative simulation trials on each hypothesis for the same phenomenon.
- A method to extract the structure of a hypothesis can be carefully designed based on a *hypothesis data representation* and shall be reducible in terms of machine-readable format for mathematical modeling, viz., W3C's **MathML**, which we shall adopt as a standard for hypothesis specification.
- We have seen that the controlled *introduction of uncertainty* into simulation data is amenable to algorithm design and then reducible to a design-theoretic synthesis method for the construction of U-relational DB's.
- Simulation data can be modeled as hypothesis data whenever it is associated with a target phenomenon. As the same phenomenon may happen to be associated with many such hypotheses, the research activity can be modeled as a *data cleaning problem* in p-DB's.

Essentially, the vision of Υ -DB comprises a design-theoretic pipeline (Fig. 1.4). For the insertion of a hypothesis k , we shall be given a **MathML**-compliant structure \mathcal{S}_k together with its simulation trial datasets \mathcal{D}_k^ℓ in raw files (e.g., .mat, .csv). Then we apply an Extract-Transform-Load (ETL) automatic procedure to generate the hypothesis 'big' fact table H_k under the trial id's.

The extracted equations are firstly encoded into fd's. Then, at any time, as many hypotheses may have been inserted into the system, the uncertainty introduction (U-intro) procedure can be applied to process the encoded fd's and synthesize the 'uncertain' U-relations that are to be eventually conditioned on observations.

Note, in Fig. 1.4, that the ETL procedure is operated in a 'local' view for each hypothesis k , while the U-intro procedure and the conditioning are operated in the 'global' view of all available hypotheses $k = 1..n$. The pipeline opens up four main tracks of technical research challenges from the ETL stage on, viz., (i) *hypothesis encoding* and (ii) *causal reasoning over fd's*, (iii) *p-DB synthesis* and (iv) *conditioning*. We address in the sequel the three first track of challenges in depth. The problem of conditioning is outlined for further work in §7.3.

Chapter 3

Hypothesis Encoding

In this chapter we address the problem of hypothesis encoding. In §3.1 we introduce notation and basic concepts of structural equations and the problem of causal ordering. In §3.2 we study the problem of extracting the causal ordering implicit in the structure of a deterministic hypothesis and show that Simon’s classical approach [24, 48] is intractable. In §3.3 then we build upon a less notorious approach of Nayak’s [49] and borrow an efficient algorithm for it that fits very well our use case for hypothesis encoding. In §3.4 we develop an encoding scheme that builds upon the idea of structural equations through an original abstraction of hypotheses ‘as data.’ In §3.5 we present experiments that attest how the encoding scheme works in practice for large hypotheses. In §3.6 we discuss related work. In §3.7 we summarize the results of this chapter.

3.1. Preliminaries: Structural Equations

Given a system of mathematical equations involving a set of variables, to build a *structural equation model* (SEM) is, essentially, to establish a one-to-one mapping between equations and variables [24]. That shall enable further detecting the hidden asymmetry between variables, i.e., their causal ordering. For instance, Einstein’s famous equation $E = mc^2$ states the equivalence of mass and energy, summarizing a theory that can be imputed two different asymmetries (for different applications), say, given a fixed amount of mass $m = m_0$ (and recall c is a constant), predict the particle’s relativistic rest energy E ; or given the particle’s rest energy, predict its mass or potential for nuclear fission.

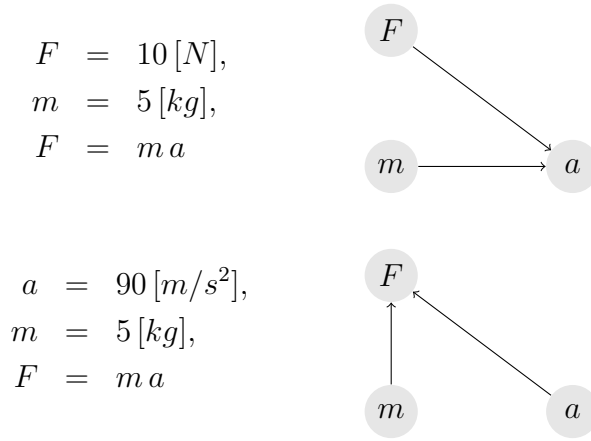


Figure 3.1. “Directed causal graphs” associated with the two systems.

To stress the point, consider Newton’s second law $F = m a$ in such a scalar setting. The modeler can either use it to compute (predict), say, acceleration values given an amount of mass and different force intensities, or to predict force intensities given a fixed acceleration (e.g., for testing an engineered dynamometer). The point here is that Newton’s equation is not enough to derive predictions. That is, it has a number of variables $|\mathcal{V}| = 3$, which is larger than $|\mathcal{E}| = 1$. It must be completed with two more equations in order to qualify as an (applied) hypothesis ‘as data.’ Although usually it is interpreted an asymmetry towards a , technically, there is nothing in its semantics to suggest so.¹ Compare the two systems given in Fig. 3.1.² In sum, the causal ordering of any system of equations is not to be guessed, as it can be inferred. In this chapter we rely on previous work (mostly AI’s work, viz., [24, 49, 48]) and adapt it for the encoding of hypotheses into fd’s.

Def. 1 A **structure** is a pair $\mathcal{S}(\mathcal{E}, \mathcal{V})$, where \mathcal{E} is a set of equations over set \mathcal{V} of variables, $|\mathcal{E}| \leq |\mathcal{V}|$, such that:

- (a) In any subset of k equations of the structure, at least k different variables appear;
- (b) In any subset of k equations in which r variables appear, $k \leq r$, if the values of any $(r - k)$ variables are chosen arbitrarily, then the values of the

¹ As the equality construct ‘=’ is used as a predicate, not an assignment operator.

² We shall introduce the notion of ‘directed causal graphs’ shortly.

remaining k variables can be determined uniquely — finding these unique values is a matter of solving the equations.

Def. 2 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a structure. We say that \mathcal{S} is self-contained or **complete** if $|\mathcal{E}| = |\mathcal{V}|$.

In short, we are interested in systems of equations that are ‘structural’ (Def. 1) and ‘complete’ (Def. 2), viz., that has as many equations as variables and no subset of equations has fewer variables than equations.³

Complete structures can be solved for unique sets of values of their variables. In this work, however, we are not concerned with solving sets of mathematical equations at all, but with processing their causal ordering in view of U-relational DB design. Simon’s concept of causal ordering has its roots in econometrics studies (cf. [24]) and has been taken further in AI with a flavor of Graphical Models (GMs) [50, 25, 48]. In this thesis we translate the problem of causal ordering into the language of data dependencies, viz., into fd’s.

Def. 3 Let \mathcal{S} be a structure. We say that \mathcal{S} is **minimal** if it is complete and there is no complete structure $\mathcal{S}' \subset \mathcal{S}$.

Def. 4 The **structure matrix** $A_{\mathcal{S}}$ of a structure $\mathcal{S}(\mathcal{E}, \mathcal{V})$, with $f_1, f_2, \dots, f_n \in \mathcal{E}$ and $x_1, x_2, \dots, x_m \in \mathcal{V}$, is a $n \times m$ matrix of 1’s and 0’s in which entry a_{ij} is non-zero if variable x_j appears in equation f_i , and zero otherwise.

Elementary row operations (e.g., row multiplication by a constant) on the structure matrix may hinder the structure’s causal ordering and then are not valid in general [24]. This also emphasizes that the problem of causal ordering is not about solving the system of mathematical equations of a structure, but identifying its hidden asymmetries.

Def. 5 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure. Then a **total causal mapping** over \mathcal{S} is a bijection $\varphi: \mathcal{E} \rightarrow \mathcal{V}$ such that, for all $f \in \mathcal{E}$, if $\varphi(f) = x$ then $x \in Vars(f)$.

³ Also, we expect the systems of equations given as input to be ‘independent’ in the sense of Linear Algebra. In our context, that means systems that can only have non-redundant equations. In that case, if some subset of equations has fewer variables than equations, then the system must be ‘overconstrained.’

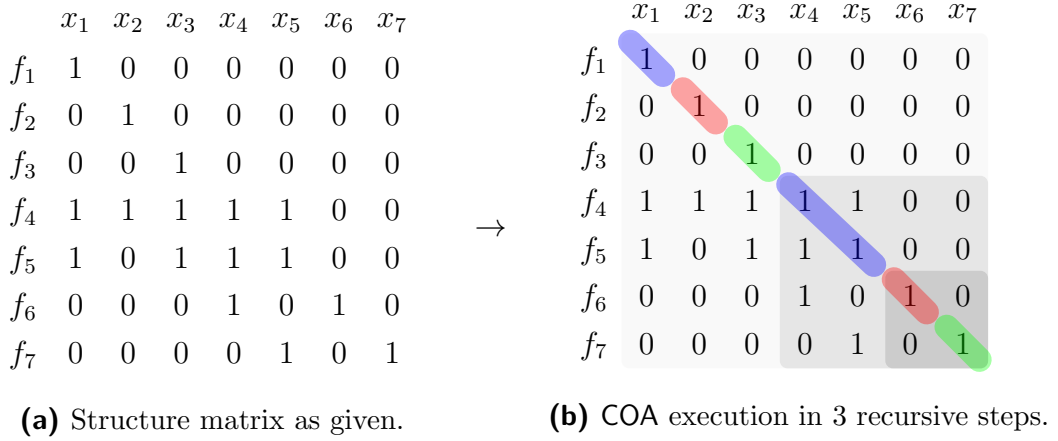


Figure 3.2. Running Simon’s Causal Ordering Algorithm (COA) on a given structure (Fig. 3.2a). Minimal subsets detected in a recursive step k , highlighted in different shades of gray, have their diagonal elements colored (Fig. 3.2b).

Simon has informally described an algorithm (cf. [24]) that, given a complete structure $\mathcal{S}(\mathcal{E}, \mathcal{V})$, can be used to compute a *partial* causal mapping φ_p from partitions on the set of equations to same-cardinality partitions on the set of variables. As shown by Dash and Druzdzal [48], the causal mapping returned by Simon’s (so-called) Causal Ordering Algorithm (COA) is not *total* when \mathcal{S} has variables that are *strongly coupled* (because they can only be determined simultaneously). They also have shown that any total mapping φ over \mathcal{S} must be consistent with COA’s partial mapping φ_p [48]. The latter is made partial by design (merge strongly coupled variables into partitions or clusters) in order to force its induced causal graph G_{φ_p} to be acyclic. Algorithm 1, COA_t , is a variant of Simon’s COA adapted to illustrate our use case. It returns a total causal mapping φ , instead of a partial causal mapping. We illustrate it through Example 3 and Fig. 3.2.

Example 3 Consider structure $\mathcal{S}(\mathcal{E}, \mathcal{V})$ whose matrix is shown in Fig. 3.2a. Note that \mathcal{S} is complete, since $|\mathcal{E}| = |\mathcal{V}| = 7$, but not minimal. The set of all minimal subsets $\mathcal{S}' \subset \mathcal{S}$ is $\mathcal{S}_c = \{\{f_1\}, \{f_2\}, \{f_3\}\}$. By eliminating the variables identified at recursive step k , a smaller structure $\mathcal{T} \subset \mathcal{S}$ is derived. Compare the partial causal mapping eventually returned by COA, $\varphi_p = \{\langle \{f_1\}, \{x_1\} \rangle, \langle \{f_2\}, \{x_2\} \rangle, \langle \{f_3\}, \{x_3\} \rangle, \langle \{f_4, f_5\}, \{x_4, x_5\} \rangle, \langle \{f_6\}, \{x_6\} \rangle, \langle \{f_7\}, \{x_7\} \rangle\}$, to the total causal mapping returned by COA_t , $\varphi = \{\langle f_1, x_1 \rangle, \langle f_2, x_2 \rangle, \langle f_3, x_3 \rangle, \langle f_4, x_4 \rangle, \langle f_5, x_5 \rangle, \langle f_6, x_6 \rangle, \langle f_7, x_7 \rangle\}$.

Algorithm 1 COA_t as a variant of Simon's COA.

```

1: procedure  $\text{COA}_t(\mathcal{S}$ : structure over  $\mathcal{E}$  and  $\mathcal{V}$ )
Require:  $\mathcal{S}$  given is complete, i.e.,  $|\mathcal{E}| = |\mathcal{V}|$ 
Ensure: Returns total causal mapping  $\varphi : \mathcal{E} \rightarrow \mathcal{V}$ 
2:    $\varphi \leftarrow \emptyset, \mathcal{S}_c \leftarrow \emptyset$ 
3:   for all minimal  $\mathcal{S}' \subset \mathcal{S}$  do
4:      $\mathcal{S}_k \leftarrow \mathcal{S}_k \cup \mathcal{S}'$  ▷ store minimal structures  $\mathcal{S}'$  found in  $\mathcal{S}$ 
5:      $\mathcal{V}' \leftarrow \mathcal{S}'(\mathcal{V})$ 
6:     for all  $f \in \mathcal{S}'(\mathcal{E})$  do
7:        $x \leftarrow \text{any } x_a \in \mathcal{V}'$ 
8:        $\varphi \leftarrow \varphi \cup \langle f, x \rangle$ 
9:        $\mathcal{V}' \leftarrow \mathcal{V}' \setminus \{x\}$ 
10:   $\mathcal{T} \leftarrow \mathcal{S} \setminus \bigcup_{\mathcal{S}' \in \mathcal{S}_k} \mathcal{S}'$ 
11:  if  $\mathcal{T} \neq \emptyset$  then
12:    return  $\varphi \cup \text{COA}_t(\mathcal{T})$ 
13:  return  $\varphi$ 

```

Since x_4 and x_5 are strongly coupled (see Fig.3.2b), COA_t maps them arbitrarily (e.g., it could be $f_4 \mapsto x_5, f_5 \mapsto x_4$ instead). Such total mapping φ renders a cycle in the directed causal graph G_φ induced by φ (see Fig.3.3). \square

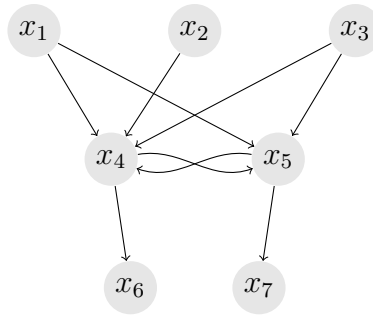


Figure 3.3. Directed causal graph G_φ induced by mapping φ for structure \mathcal{S} . An edge connects a node x_i towards a node x_j , with $x_i, x_j \in \mathcal{V}$, iff x_i appears in the equation $f \in \mathcal{E}$ such that $\varphi(f) = x_j$.

3.2. The Problem of Causal Ordering

The serious issue with Alg. 1, COA_t , is that finding all minimal structures in a given structure (cf. line 3) is a hard problem that can only be addressed heuristically as a problem of co-clustering (also called biclustering [51, 52]) in

Boolean matrices. Simon’s approach, however, as we shall see next, is not the only way to cope with the problem of causal ordering.

In fact, in order to study the computational properties of SEM’s and the problem of causal ordering, we observe that any structure $\mathcal{S}(\mathcal{E}, \mathcal{V})$ satisfying Def. 1 can be modeled straightforwardly as a bipartite graph $G = (V_1 \cup V_2, E)$, where the set \mathcal{E} of equations and the set \mathcal{V} of variables are the disjoint vertex sets, i.e., $V_1 \mapsto \mathcal{E}$, $V_2 \mapsto \mathcal{V}$, and $E \mapsto \mathcal{S}$ is the edge set connecting equations to the variables appearing in them. Fig. 3.4 shows the bipartite graph G corresponding to the structure given in Example 3 — for a comprehensive text on graph concepts and its related algorithmic problems, cf. Even [53].

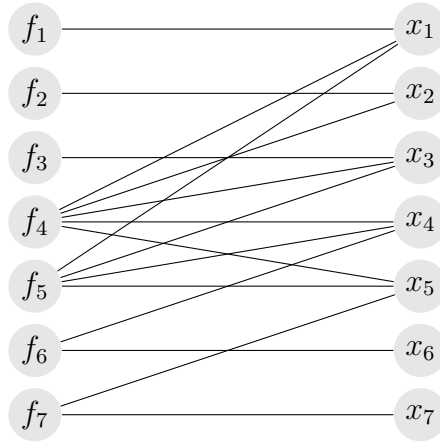


Figure 3.4. Bipartite graph G of structure \mathcal{S} from Example 3.

A *biclique* (or complete bipartite graph) is a bipartite graph $G = (A \cup B, E)$ such that for every two vertices $a \in A$, $b \in B$, we have $(a, b) \in E$ [53]. Note that for *balanced* bicliques, i.e., when $|A| = |B| = K$, the degree $\deg(u)$ of any vertex $u \in A \cup B$ must be $\deg(u) = |A| = |B| = K$.

Recent approaches to co-clustering problems (e.g., [54]) have come with the notion of *pseudo-biclique* (also called ‘quasi-biclique’), which is a relaxation of the biclique concept to allow some less rigid notion of connectivity than the ‘complete connectivity’ required in a biclique. Now, recall that Simon’s COA_t needs to find, at each recursive step, all minimal subsets $\mathcal{S}' \subseteq \mathcal{S}$. Theorem 1 situates this particular computational task in terms of its complexity, which for $|\mathcal{E}'| = |\mathcal{V}'| \geq 2$ is equivalent to find, at each recursive step, the minimal-size ‘pseudo-bicliques’

(i.e., with the least $K \geq 2$) in its corresponding bipartite graph (e.g., see Fig. 3.4). Here we take, in Def. 6, a specific notion of pseudo-biclique.

Def. 6 Let $G = (A \cup B, E)$ be a bipartite graph. We say that G is a **K -balanced pseudo-biclique** if $|A| = |B| = K$ with $|E| \geq 2K$ and, for all vertices $u \in A \cup B$, $\deg(u) \geq 2$.

Now we state (originally) the *balanced pseudo-biclique problem* (BPBP) as a decision problem as follows.

(BPBP). Given a bipartite graph $G = (V_1 \cup V_2, E)$ and a positive integer $K \geq 2$, does G contain a K -balanced pseudo-biclique?

Lemma 1 The *balanced pseudo-biclique problem* (BPBP) is NP-Complete.

Proof 1 We show (by restriction [55]) that the BPBP is a generalization of the balanced biclique problem (BBP), referred ‘balanced complete bipartite subgraph’ problem [55, GT24, p. 196], which is shown to be NP-Complete by means of a transformation from ‘clique’ [56, p. 446]. The restriction from BPBP to BBP (special case) is made by requiring (cf. Def. 6) either (a) $|E| = K^2$ or (b) $\deg(u) = K$, for $K \geq 2$,⁴ which are equivalent ways of enforcing the inquired K -balanced pseudo-biclique to be a K -balanced biclique. \square

We introduce another hypothesis structure (see Fig. 3.5) to illustrate the correspondence between the pseudo-biclique property and COA’s algorithmic approach as elaborated in the proof of Theorem 1.

Theorem 1 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure. Then the extraction of its causal ordering by Simon’s $\text{COA}(\mathcal{S})$ is NP-Hard.

Proof 2 We show that, at each recursive step k of COA, to find all non-trivial minimal subsets (i.e., $|\mathcal{E}'| \geq 2$) translates into an optimization problem associated with the decision problem BPBP, which we know by Lemma 1 to be NP-Complete. See Appendix §A.1.1. \square

⁴ Note that clearly, for any positive integer $K \geq 2$, we have for $K^2 \geq 2K$.

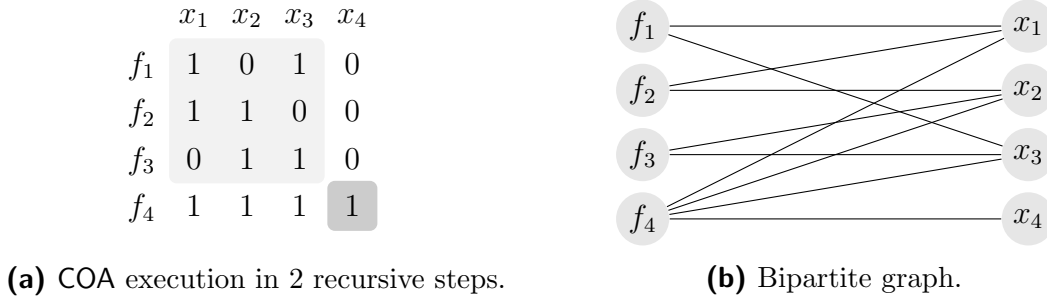


Figure 3.5. Another hypothesis structure example.

Nonetheless, the problem of causal ordering can be solved efficiently by means of a different, less notorious approach due to Nayak [49], which we introduce and build upon next.

3.3. Total Causal Mappings

The problem of causal ordering can be solved in polynomial time by (i) finding any total causal mapping $\varphi: \mathcal{E} \rightarrow \mathcal{V}$ over structure \mathcal{S} given (cf. Def. 5); and then (ii) by computing the transitive closure C_φ^+ of the set C_φ (cf. Eq. 3.1) of *direct causal dependencies* induced by φ .

$$C_\varphi = \{ (x_a, x_b) \mid \text{there exists } f \in \mathcal{E} \text{ such that } \varphi(f) = x_b \text{ and } x_a \in \text{Vars}(f) \} \quad (3.1)$$

Def. 7 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a structure with variables $x_a, x_b \in \mathcal{V}$, and φ a total causal mapping over \mathcal{S} inducing set of direct causal dependencies C_φ and its transitive closure C_φ^+ . We say that (x_b, x_a) is a **direct causal dependency** in \mathcal{S} if $(x_b, x_a) \in C_\varphi$, and that (x_b, x_a) is a **causal dependency** in \mathcal{S} if $(x_b, x_a) \in C_\varphi^+$.

In other words, (x_a, x_b) is in C_φ iff x_b direct and causally depends on x_a , given the causal asymmetries induced by φ . Those notions open up an approach to causal reasoning that fits very well to our use case, which is aimed at encoding hypothesis structures into fd sets and then performing (symbolic) causal reasoning in terms of acyclic pseudo-transitive reasoning over fd's (cf. Chapter 4).⁵ For it to be effective, nonetheless, we shall need to ensure some properties of total causal mappings first.

⁵ Note that it differs from AI research (e.g., [48]) geared for reasoning over GM's.

For a given structure \mathcal{S} , there may be multiple total causal mappings over \mathcal{S} (recall Example 3). But the causal ordering of \mathcal{S} must be unique (see Fig. 3.3). Therefore, a question that arises is whether the transitive closure C_φ^+ is the same for any total causal mapping φ over \mathcal{S} . Proposition 1, originally from Nayak [49], ensures that is the case.

Proposition 1 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a structure, and $\varphi_1: \mathcal{E} \rightarrow \mathcal{V}$ and $\varphi_2: \mathcal{E} \rightarrow \mathcal{V}$ be any two total causal mappings over \mathcal{S} . Then $C_{\varphi_1}^+ = C_{\varphi_2}^+$.

Proof 3 The proof is based on an argument from Nayak [49], which we present in arguably much clearer way (see Appendix, §A.1.2). Intuitively, it shows that if φ_1 and φ_2 differ on the variable an equation f is mapped to, then such variables, viz., $\varphi_1(f)$ and $\varphi_2(f)$, must be causally dependent on each other (strongly coupled). \square

Another issue is concerned with the precise conditions under which total causal mappings exist (i.e., whether or not all variables in the equations can be causally determined). In fact, by Proposition 2, based on Nayak [49] apud. Hall [53, p. 135-7], we know that the existence condition holds iff the given structure is complete.

Before proceeding to it, let us refer to Even [53] to briefly introduce the additional graph-theoretic concepts which are necessary here. A *matching* in a graph is a subset of edges such that no two edges in the matching share a common node. A matching is said *maximum* if no edge can be added to the matching (without hindering the matching property). Finally, a matching in a graph is said ‘perfect’ if every vertex is an end-point of some edge in the matching — in a bipartite graph, a perfect matching is said a *complete* matching.

Proposition 2 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a structure. Then a total causal mapping $\varphi: \mathcal{E} \rightarrow \mathcal{V}$ over \mathcal{S} exists iff \mathcal{S} is complete.

Proof 4 We observe that a total causal mapping $\varphi: \mathcal{E} \rightarrow \mathcal{V}$ over \mathcal{S} corresponds exactly to a complete matching M in a bipartite graph $B = (V_1 \cup V_2, E)$, where $V_1 \mapsto \mathcal{E}$, $V_2 \mapsto \mathcal{V}$, and $E \mapsto \mathcal{S}$. In fact, by Even apud. Hall’s theorem (cf. [53, 135-7]), we know that B has a complete matching iff (a) for every subset of vertices

$F \subseteq V_1$, we have $|F| \leq |E(F)|$, where $E(F)$ is the set of all vertices connected to the vertices in F by edges in E ; and (b) $|V_1| = |V_2|$. By Def. 1 (no subset of equations has fewer variables than equations), and Def. 2 (number of equations is the same as number of variables), it is easy to see that conditions (a) and (b) above hold iff \mathcal{S} is a complete structure. \square

The problem of finding a maximum matching is a well-studied algorithmic problem. In this thesis we adopt the Hopcroft-Karp algorithm [57], which is known to be polynomial-time, bounded by $O(\sqrt{|V_1| + |V_2|} |E|)$.⁶ That is, we handle the problem of total causal mapping by (see Alg. 2) translating it to the problem of maximum matching in a bipartite graph (in linear time) and then applying the Hopcroft-Karp algorithm to get the matching and finally translate it back to the total causal mapping, as suggested by the proof of Proposition 2.

Algorithm 2 Find a total causal mapping for a given structure.

```

1: procedure TCM( $\mathcal{S}$ : structure over  $\mathcal{E}$  and  $\mathcal{V}$ )
Require:  $\mathcal{S}$  given is a complete structure, i.e.,  $|\mathcal{E}| = |\mathcal{V}|$ 
Ensure: Returns a total causal mapping  $\varphi$ 
2:    $B(V_1 \cup V_2, E) \leftarrow \emptyset$ 
3:    $\varphi \leftarrow \emptyset$ 
4:   for all  $\langle f, X \rangle \in \mathcal{S}$  do  $\triangleright$  translates the structure  $\mathcal{S}$  to a bipartite graph  $B$ 
5:      $V_1 \leftarrow V_1 \cup \{f\}$ 
6:     for all  $x \in X$  do
7:        $V_2 \leftarrow V_2 \cup \{x\}$ 
8:        $E \leftarrow E \cup \{(f, x)\}$ 
9:    $M \leftarrow \text{Hopcroft-Karp}(B)$   $\triangleright$  solves the maximum matching problem
10:  for all  $(f, x) \in M$  do  $\triangleright$  translates the matching to a total causal mapping
11:     $\varphi \leftarrow \varphi \cup \{\langle f, x \rangle\}$ 
12:  return  $\varphi$ 

```

Fig. 3.6 shows the complete matching found by the Hopcroft-Karp algorithm for the structure given in Example 3.

⁶ The Hopcroft-Karp algorithm solves maximum matching in a bipartite graph efficiently as a problem of finding maximum flow in a network (cf. [53, p. 135-7], or [58, p. 664-9]).

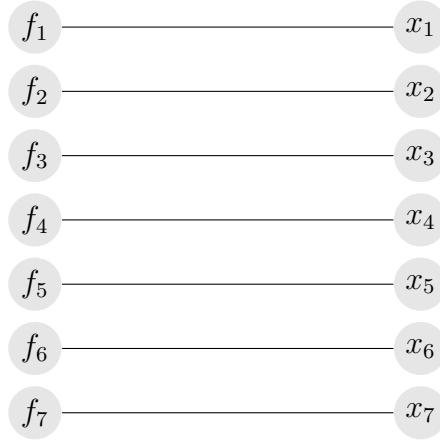


Figure 3.6. Complete matching M for structure S from Example 3.

Corollary 1 summarizes the results we have so far.

Corollary 1 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure. Then a total causal mapping $\varphi: \mathcal{E} \rightarrow \mathcal{V}$ over \mathcal{S} can be found by (Alg. 2) TCM in time that is bounded by $O(\sqrt{|\mathcal{E}|} |\mathcal{S}|)$.

Proof 5 Let $B = (V_1 \cup V_2, E)$ be the bipartite graph corresponding to complete structure \mathcal{S} given to TCM, where $V_1 \mapsto \mathcal{E}$, $V_2 \mapsto \mathcal{V}$, and $E \mapsto \mathcal{S}$. The translation of \mathcal{S} into B is done by a scan over it. This scan is of length $|\mathcal{S}| = |E|$. Note that number $|E|$ of edges rendered is precisely the length $|\mathcal{S}|$ of structure, where the denser the structure, the greater $|\mathcal{S}|$ is. The re-translation of the matching computed by internal procedure Hopcroft-Karp, in turn, is done at expense of $|\mathcal{E}| = |\mathcal{V}| \leq |\mathcal{S}|$. Thus, it is easy to see that TCM is dominated by the maximum matching algorithm Hopcroft-Karp, which is known to be $O(\sqrt{|V_1| + |V_2|} |E|)$, i.e., $O(\sqrt{|\mathcal{E}| + |\mathcal{V}|} |\mathcal{S}|)$. Since \mathcal{S} is assumed complete, we have $|\mathcal{E}| = |\mathcal{V}|$ then $\sqrt{|\mathcal{E}| + |\mathcal{V}|} = \sqrt{2} \sqrt{|\mathcal{E}|}$. Therefore, TCM must have running time at most $O(\sqrt{|\mathcal{E}|} |\mathcal{S}|)$. \square

Remark 2 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure. Then we know (cf. Proposition 2) that a total causal mapping over \mathcal{S} exists. Let it be defined $\varphi \triangleq \text{TCM}(\mathcal{S})$. Then the causal ordering implicit in \mathcal{S} shall be correctly extracted (cf. Proposition 1) by processing the causal dependencies induced by φ , as we show in Chapter 4. \square

Now we are ready to accomplish the hypothesis encoding into fd's, as we show next.

3.4. The Encoding Scheme

We shall encode variables as relational attributes and map equations onto fd's through total causal mappings. Let Z be a set of attribute symbols such that $Z \simeq \mathcal{V}$, where $\mathcal{S}(\mathcal{E}, \mathcal{V})$ is a complete structure; and let $\phi, v \notin Z$ be two special attribute symbols kept to identify (resp.) phenomena and hypotheses. We are explicitly distinguishing symbols in Z , assigned by the user into structure \mathcal{S} , from epistemological symbols ϕ and v . Then we consider a sense of Simon's into the nature of scientific modeling and interventions [24], summarized in Def. 8.

Def. 8 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a structure and $x \in \mathcal{V}$ be a variable. We say that x is **exogenous** if there exists an equation $f \in \mathcal{E}$ such that $\text{Vars}(f) = \{x\}$. In this case f can be written $f(x) = 0$, and must be mapped to x in any total causal mapping φ over \mathcal{S} . We say that x is **endogenous** otherwise.

Remark 3 introduces an interpretation of Def. 8 with a data dependency flavor.

Remark 3 The values of exogenous variables (attributes) are to be determined empirically, outside of the system (proposed structure \mathcal{S}). Such values are, therefore, dependent on the phenomenon id ϕ only. The values of endogenous variables (attributes) are in turn to be determined theoretically, within the system. They are dependent on the hypothesis id v and shall be dependent on the phenomenon id ϕ as well indirectly. \square

As introduced in §2.2, the encoding scheme we are presenting here is not obvious. It goes beyond Simon's structural equations to abstract the data-level semantics of mathematical deterministic hypotheses. Whereas Simon's structural equations are able to represent only linear equations, our encoding scheme can represent non-linear equations and arbitrarily complex mathematical operators by means of its data representation of deterministic hypotheses.

For instance, take non-linear equation $y = a x^2$ and suppose that, considering the context of its complete system of equations, (Alg. 2) TCM maps it onto variable y . Then, by an abstraction of the equation semantics, we shall encode it into fd

$$\Sigma = \{ \begin{array}{l} \phi \rightarrow x_1, \\ \phi \rightarrow x_2, \\ \phi \rightarrow x_3, \\ x_1 x_2 x_3 x_5 v \rightarrow x_4, \\ x_1 x_3 x_4 v \rightarrow x_5, \\ x_4 v \rightarrow x_6, \\ x_5 v \rightarrow x_7 \end{array} \}.$$

Figure 3.7. Encoded fd set Σ (cf. Alg. 3) for the structure from Example 3.

$axv \rightarrow y$. That is, the hypothesis identifier v captures the data-level semantics of the hypothesis equation.⁷

We encode complete structures into fd sets by means of (Alg. 3) **h-encode**. Fig. 4.1 presents an fd set defined $\Sigma \triangleq \text{h-encode}(\mathcal{S})$, encoding the same structure \mathcal{S} from Example 3.

Algorithm 3 Hypothesis encoding.

```

1: procedure H-ENCODE( $\mathcal{S}$ : structure over  $\mathcal{E}$  and  $\mathcal{V}$ ,  $\mathcal{D}$ : domain variables)
Require:  $\mathcal{S}$  given is a complete structure, i.e.,  $|\mathcal{E}| = |\mathcal{V}|$ 
Ensure: Returns a non-redundant fd set  $\Sigma$ 
2:    $\Sigma \leftarrow \emptyset$ 
3:    $\varphi \leftarrow \text{TCM}(\mathcal{S})$ 
4:   for all  $\langle f, x \rangle \in \varphi_t$  do
5:      $Z \leftarrow X \setminus \{x\}$ , where  $\langle f, X \rangle \in \mathcal{S}$ 
6:     if  $Z = \emptyset$  or  $Z \subseteq \mathcal{D}$  then  $\triangleright x$  is exogenous
7:       if  $x \notin \mathcal{D}$  then  $\triangleright$  suppress  $\phi$ -fd for dimensions like time  $t$ 
8:          $\Sigma \leftarrow \Sigma \cup \langle Z \cup \{\phi\}, x \rangle$ 
9:       else  $\triangleright x$  is endogenous
10:         $\Sigma \leftarrow \Sigma \cup \langle Z \cup \{v\}, x \rangle$ 
11:   return  $\Sigma$ 

```

Now we study the design-theoretic properties of the encoded fd sets. We shall make use of the concept of ‘canonical’ fd sets (also called ‘minimal’ [20, p. 390]), see Def. 9.

Def. 9 Let Σ be an fd set. We say that Σ is **canonical** if:

⁷ Note that, without the hypothesis id, infinitely many equations fit the pattern $ax \rightarrow y$.

- (a) each fd in Σ has the form $X \rightarrow A$, where $|A| = 1$;
- (b) For no $\langle X, A \rangle \in \Sigma$ we have $(\Sigma - \{\langle X, A \rangle\})^+ = \Sigma^+$;
- (c) for each fd $X \rightarrow A$ in Σ , there is no $Y \subset X$ such that $(\Sigma \setminus \{X \rightarrow A\} \cup \{Y \rightarrow A\})^+ = \Sigma^+$.

For an fd set satisfying such properties (Def. 9) individually, we say that it is (a) *singleton-rhs*, (b) *non-redundant* and (c) *left-reduced*. It is said to have an attribute A in X that is ‘extraneous’ w.r.t. Σ if it is *not* left-reduced (Def. 9-c) [22, p. 74]. Finally, an fd $X \rightarrow Y$ in Σ is said *trivial* if $Y \subseteq X$. Note that the presence of a trivial fd in a an fd set is sufficient to make it redundant.

Theorem 2 Let Σ be an fd set defined $\Sigma \triangleq \text{h-encode}(\mathcal{S})$ for some complete structure \mathcal{S} . Then Σ is non-redundant and singleton-rhs but may not be left-reduced (then may not be canonical).

Proof 6 We show that properties (a-b) of Def. 9 must hold for Σ produced by (Alg. 3) *h-encode*, but property (c) may not hold (i.e., encoded fd set Σ may not be left-reduced). See Appendix, §A.1.3. \square

We draw attention to the significance of Theorem 2, as it sheds light on a connection between Simon’s complete structures [24] and fd sets [20]. In fact, we continue to elaborate on that connection in next chapter to handle causal ordering processing symbolically by causal reasoning over fd’s.

3.5. Experiments

Fig. 3.8 shows the results of experiments we have carried out in order to study how effective the procedure of hypothesis encoding is in practice, in particular its behavior for hypotheses whose structure \mathcal{S} has been randomly generated over orders of magnitude $|\mathcal{S}| \approx 2^k$, to have length up to $|\mathcal{S}| \approx 2^{20} \lesssim 1M$. The largest structure considered, with $|\mathcal{S}| \approx 1M$, has been generated to have exactly $|\mathcal{E}| = 2.5K$.

We have executed ten runs for each tested order of magnitude, and then taken its mean running time in *ms*.⁸ The plot is shown in Fig. 3.8 in logscale base 2. In fact a near-, sub-quadratic slope is expected for the curve structure length $|\mathcal{S}| \times \text{time}$.

These scalability results are compatible with the computational complexity of *h-encode*, which is (cf. Corollary 1) bounded by $O(\sqrt{|\mathcal{E}|} |\mathcal{S}|)$.⁹

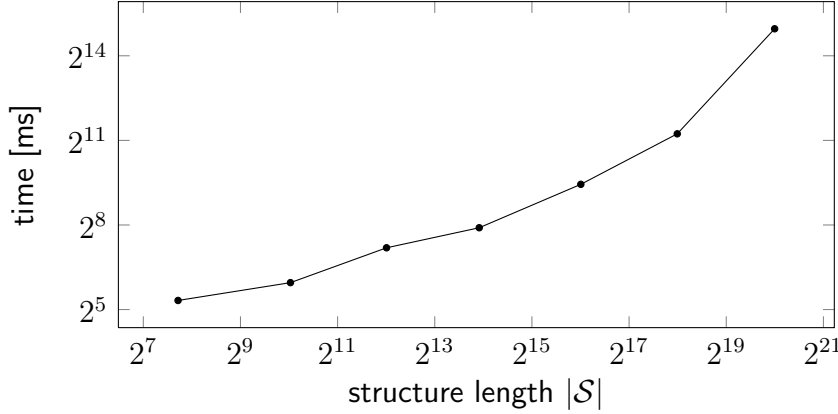


Figure 3.8. Performance of hypothesis encoding (in logscale).

3.6. Related Work

Modeling physical and socio-economical systems as a set of equations is a traditional modeling approach, and a very large bulk of models exist up to date. Simon's early work on structural equations and causal ordering comprises a specific notion of causality aimed at further contributing to the potential of such modeling approach (cf. [59]). It is meant for *identifying* influences among variables (or their values) that are implicit in the system model for enabling *informed* interventions. These may apply either to the system (phenomenon) under study, or to the model itself (say, when its predictions are not approximating observations very well).

Significant research effort has been devoted to causal modeling and reasoning in the past decades in both statistics and AI (cf. [25, 19]). The notion of causality used can be traced back to the early work of Simon's and others in Econometrics.

⁸ The experiments were performed on a 2.3GHZ/4GB Intel Core i5 running Mac OS X 10.6.8.

⁹ Note that, for any arbitrary structure $\mathcal{S}(\mathcal{E}, \mathcal{V})$, we have $|\mathcal{E}| \leq |\mathcal{S}| \leq |\mathcal{E}|^2$. So, in worst case (the densest structure possible) we have $|\mathcal{S}| = |\mathcal{E}|^2$ and then can establish a time bound function of $|\mathcal{E}|$ only, viz., $O(|\mathcal{E}|^2 \sqrt{|\mathcal{E}|})$.

Nonetheless, there are two important differences to be emphasized:

- Such work is majorly devoted to deal with (statistical) qualitative hypotheses, not (deterministic) quantitative hypotheses;
- The causal model is assumed as given or is derived from data, instead of being converted or synthesized from a set of equations.

These are both core differences that also apply to our work in comparison to the bulk of existing work in probabilistic DB's. Our main point here, though, is to clarify the technical context and state of the art of the problem of causal ordering. A few works have been concerned with extracting a causal model out of some previous existing formal specification such as a set of equations. This is a reason why causal ordering has been an yet barely studied problem from the computational point of view.

Dash and Druzdzel revisit the problem and re-motivate it in light of modern applications [48]. First, they provide a formal description of how Simon's COA gives a summary of the causal dependencies implicit in a given SEM. That is, in clustering the strongly coupled variables into a causal graph, COA provides a condensed representation of the causal model implicit in the given SEM. They show then that any valid total causal mapping produced for a given SEM must be consistent with COA's partial causal mapping.

Yet, the serious problem is that the algorithm turns out to be intractable. In fact, no formal study of COA's computational properties can yet be found in the literature. In this thesis we have obtained the (negative) hardness result that it is intractable, which turns out to be compatible with Nayak's intuition (sic.) that it is a worst-case exponential time algorithm (cf. [60, p. 37]).

Inspired on Serrano and Gossard's work on constraint modeling and reasoning [61], Nayak reports an approach that is provably quite effective to process the causal ordering: extract a total causal mapping and then compute the transitive closure of the direct causal dependencies. In this thesis we build upon it to perform causal reasoning in terms of a form of transitive reasoning. Such approach fits very well to our use case, viz., the synthesis of p-DB's from fd's. As we show

in Chapter 4, we process the causal ordering of a hypothesis structure (abstracted as a SEM) in terms of acyclic causal reasoning over fd's and prove its correctness. This is enabled by the encoding scheme presented in this chapter.

3.7. Summary of Results

In this chapter we have studied and developed an encoding scheme to process the mathematical structure of a deterministic hypothesis into a set of fd's towards the encoding of hypotheses 'as data.' Then we have studied the design-theoretic properties held by such an encoded fd set. We list the results achieved as follows.

- By Theorem 1, we know (an original hardness result) that Simon's approach to process the causal ordering of a structure is intractable;
- By building upon on the work of Simon [24] and Nayak [49] (cf. Propositions 1 and 2), we have framed an approach to efficiently extract the basic information (a total causal mapping) for processing the causal ordering implicit in the mathematical structure of a deterministic hypothesis;
- By Corollary 1, we know how to process the complete structure of a hypothesis into a total causal mapping in time that is bounded by $O(\sqrt{|\mathcal{E}|} |\mathcal{S}|)$. That is, the machinery of hypothesis encoding is provably suitable for very large hypothesis structures.
- By Theorem 2, which studies the design-theoretic properties of the encoded fd sets, we have unraveled the connection between Simon's complete structures and fd sets to further explore it in next chapter.
- We have performed experiments (cf. Fig. 3.8) to study how effective the approach is in practice, or how it scales for hypotheses whose structure \mathcal{S} is randomly generated to have length up to the order of $|\mathcal{S}| \lesssim 1M$.¹⁰

¹⁰ The tests were up to this order only because of the hardware limitations of our experimental settings. In theory (cf. complexity time bounds), larger structures can be handled very efficiently.

Chapter 4

Causal Reasoning over FD's

In this chapter we present a technique to address the problem of causal ordering processing in order to enable the synthesis of U-relational DB's. In §4.1 we introduce Armstrong's classical inference system to reason over fd's. In §4.2 we develop the core concept and algorithm of the folding of an fd set, as a method for acyclic causal reasoning over fd's. In §4.3 we show its connections (equivalence) with causal reasoning. In §4.4 we present experiments on how the method behaves in practice. §4.5 we discuss related work. In §4.6 we conclude the chapter.

4.1. Preliminaries: Armstrong's Inference Rules

As usual notational conventions from the DB literature [20, 21], we write X, Y, Z to denote sets of relational attributes and A, B, C to denote singleton attribute sets. Also, we write XY as shorthand for $X \cup Y$.

Functional dependency theory relies on Armstrong's inference rules (or axioms) of (R0) reflexivity, (R1) augmentation and (R2) transitivity, which forms a sound and complete inference system for reasoning over fd's [20]. From R0-R2 one can derive additional rules, viz., (R3) decomposition, (R4) union and (R5) pseudo-transitivity.

- R0.** If $Y \subseteq X$, then $X \rightarrow Y$;
- R1.** If $X \rightarrow Y$, then $XZ \rightarrow YZ$;
- R2.** If $X \rightarrow Y$ and $Y \rightarrow W$, then $X \rightarrow W$;
- R3.** If $X \rightarrow YZ$, then $X \rightarrow Y$ and $X \rightarrow Z$;
- R4.** If $X \rightarrow Y$ and $X \rightarrow Z$, then $X \rightarrow YZ$;

R5. If $X \rightarrow Y$ and $YZ \rightarrow W$, then $XZ \rightarrow W$.

Given an fd set Σ , one can obtain Σ^+ , the closure of Σ , by a finite application of rules R0-R5. We are concerned with reasoning over an fd set in order to process its implicit causal ordering. The latter, as we shall see in §4.3, can be performed in terms of (pseudo-)transitive reasoning. Note that R2 is a particular case of R5 when $Z = \emptyset$, then we shall refer to R5 reasoning and understand R2 included. The next definition opens up a way to compute Σ^+ very efficiently.

Let Σ be an fd set on attributes U , with $X \subseteq U$. Then X^+ , the *attribute closure* of X w.r.t. Σ , is the set of attributes A such that $\langle X, A \rangle \in \Sigma^+$. Bernstein has long given algorithm (Alg. 4) **XClosure** to compute X^+ . It is polynomial time in $|\Sigma| \cdot |U|$ (cf. [62]), where Σ and U are (resp.) the given fd set and the attribute set over which it is defined. A tighter time bound (linear time in $|\Sigma| \cdot |U|$) is achievable as discussed further in Remark 4.

Algorithm 4 Attribute closure X^+ (cf. [20, p. 388]).

```

1: procedure XCLOSURE( $\Sigma$ : fd set,  $X$ : attribute set)
Require:  $\Sigma$  is an fd set,  $X$  is a non-empty attribute set
Ensure:  $X^+$  is the attribute closure of  $X$  w.r.t.  $\Sigma$ 
2:   size  $\leftarrow 0$ 
3:    $\Delta \leftarrow \emptyset$ 
4:    $X^+ \leftarrow X$ 
5:   while size <  $|X^+|$  do
6:     size  $\leftarrow |X^+|$ 
7:      $\Sigma \leftarrow \Sigma \setminus \Delta$ 
8:     for all  $\langle Y, Z \rangle \in \Sigma$  do
9:       if  $Y \subseteq X^+$  then
10:          $\Delta \leftarrow \Delta \cup \{\langle Y, Z \rangle\}$  ▷ consumes fd
11:          $X^+ \leftarrow X^+ \cup Z$ 
12:   return  $X^+$ 

```

4.2. Acyclic Pseudo-Transitive Reasoning

As discussed in the previous chapter, we shall process the causal ordering in terms of computing the transitive closure of each endogenous variable (predictive attribute). Before we proceed to that, we shall develop some machinery to reason over fd's in terms of Armstrong's rule R5 (pseudo-transitivity). We shall then

demonstrate the correspondence between this kind of reasoning with causal reasoning shortly in the sequel.

Def. 10 Let Σ be a set of fd's on attributes U . Then Σ^\triangleright , the **pseudo-transitive closure** of Σ , is the minimal set $\Sigma^\triangleright \supseteq \Sigma$ such that $X \rightarrow Y$ is in Σ^\triangleright , with $XY \subseteq U$, iff it can be derived from a finite (possibly empty) application of rule R5 over fd's in Σ . In that case, we may write $X \xrightarrow{\triangleright} Y$ and omit 'w.r.t. Σ ' if it can be understood from the context.

We are in fact interested in a very specific proper subset of Σ^\triangleright , say, a kernel of fd's in Σ^\triangleright that gives a “compact” representation of the causal ordering implicit in Σ . Note that, to characterize such special subset we shall need to be careful w.r.t. the presence of cycles in the causal ordering.

Def. 11 Let Σ be a set of fd's on attributes U , and $\langle X, A \rangle \in \Sigma^\triangleright$ with $XA \subseteq U$. We say that $X \rightarrow A$ is **folded** (w.r.t. Σ), and write $X \xrightarrow{\triangleright} A$, if it is non-trivial and for no $Y \subset U$ with $Y \not\supseteq X$, we have $Y \rightarrow X$ and $X \not\rightarrow Y$ in Σ^+ .

The intuition of Def. 11 is that an fd is folded when there is no sense in going on with pseudo-transitive reasoning over it anymore (nothing new is discovered). Given an fd $X \rightarrow A$ in fd set Σ , we shall be able to find some folded fd $Z \rightarrow A$ by applying (R5) pseudo-transitivity as much as possible while ruling out cyclic or trivial fd's in some clever way.

Def. 12 Let Σ be an fd set on attributes U , and $\langle X, A \rangle \in \Sigma$ be an fd with $XA \subseteq U$. Then,

- (a) A^{\triangleright} , the **(attribute) folding** of A (w.r.t. Σ) is an attribute set $Z \subset U$ such that $Z \xrightarrow{\triangleright} A$;
- (b) Accordingly, Σ^{\triangleright} , the **folding** of Σ , is a proper subset $\Sigma^{\triangleright} \subset \Sigma^\triangleright$ such that an fd $\langle Z, A \rangle \in \Sigma^\triangleright$ is in Σ^{\triangleright} iff $X \xrightarrow{\triangleright} A$ for some $Z \subset U$.

Example 4 (continued). Fig. 4.1 shows an fd set Σ (left) and its folding Σ^{\triangleright} (right). Note that the folding can be obtained by computing the attribute folding

$\Sigma = \{$	$\phi \rightarrow x_1,$	$\Sigma^{\text{q}\dagger} = \{$	$\phi \rightarrow x_1,$
	$\phi \rightarrow x_2,$		$\phi \rightarrow x_2,$
	$\phi \rightarrow x_3,$		$\phi \rightarrow x_3,$
$x_1 x_2 x_3 x_5 v \rightarrow x_4,$		$\phi v x_5 \rightarrow x_4,$	
$x_1 x_3 x_4 v \rightarrow x_5,$		$\phi v x_4 \rightarrow x_5,$	
$x_4 v \rightarrow x_6,$		$\phi v x_5 \rightarrow x_6,$	
$x_5 v \rightarrow x_7 \}$		$\phi v x_4 \rightarrow x_7 \}$	

Figure 4.1. Fd set Σ encoding (cf. Alg. 3) the structure of Fig. 3.2a and its folding $\Sigma^{\text{q}\dagger}$ derived by Alg. 5.

for A in each fd $X \rightarrow A$ in Σ . We illustrate below some reasoning steps to partially compute an attribute folding considering the subset of fd's in Σ with $\phi \notin X$.

1. $x_1 x_2 x_3 x_5 v \rightarrow x_4$ [given]
 2. $x_1 x_3 x_4 v \rightarrow x_5$ [given]
 3. $x_5 v \rightarrow x_7$ [given]
 4. $x_1 x_3 x_4 v \rightarrow x_7$ [R5 over (2), (3)]
 5. $x_1 x_2 x_3 x_5 v \rightarrow x_7$ [R5 over (1), (4)]
-
6. $\therefore x_1 x_2 x_3 x_4 v \rightarrow x_7$ [R5 over (2), (5)].

Note that (6) is still amenable to further application of R5, say over (1), to derive (7) $x_1 x_2 x_3 x_5 v \rightarrow x_7$. However, even though (1) and (6) have (resp.) the form (1) $X \rightarrow A$ and (6) $Y \rightarrow B$ with $Y \xrightarrow{\text{p}} X$, we have $X \xrightarrow{\text{p}} Y$ as well which characterizes a cycle that fetches nothing into Y .¹ In fact, if we consider only the fd's 1-3 given, then (6) itself satisfies Def. 11 and then is folded. The same holds, e.g., for (1) by an empty application of R5. \square

Lemma 2 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, φ a total causal mapping over \mathcal{S} and Σ an fd set encoded through φ given \mathcal{S} . If $\langle X, A \rangle \in \Sigma$, then $A^{\text{q}\dagger}$, the attribute folding of A (w.r.t. Σ) exists and is unique.

Proof 7 See Appendix, §A.2.1. \square

¹ Note, when at the step deriving (6) by R5 over (2), (5), that such cycle was not yet formed.

We give an original algorithm (Alg. 5) to compute the folding of an fd set. At its core there lies (Alg. 6) **AFolding**, which can be understood as a non-obvious variant of **XClosure** (cf. Alg. 4) designed for acyclic pseudo-transitivity reasoning. In order to compute the folding of attribute A in fd $\langle X, A \rangle \in \Sigma$, algorithm **AFolding** can be seen as backtracing the causal ordering implicit in Σ towards A . Analogously, in terms of the directed graph G_φ induced by the causal ordering (see Fig. 3.3), that would comprise graph traversal to identify the nodes $x \in \mathcal{V}$ that have $x_a \in \mathcal{V}$ in their reachability, i.e., $x \rightsquigarrow x_a$. Rather, **AFolding**'s processing of the causal ordering is fully symbolic based on Armstrong's rewrite rule R5.

Example 5 Cyclicity in an fd set Σ may have the effect of making its folding Σ^{q} to degenerate to Σ itself. For instance, consider $\Sigma = \{A \rightarrow B, B \rightarrow A\}$. Note that Σ is canonical, and **AFolding** (w.r.t. Σ) is B given A , and A given B . That is, $\Sigma^{\text{q}} = \Sigma$. \square

Algorithm 5 Folding of an fd set.

1: **procedure** FOLDING(Σ : fd set)

Require: Σ given encodes complete structure \mathcal{S}

Ensure: Returns fd set Σ^{q} , the folding of Σ

2: $\Sigma^{\text{q}} \leftarrow \emptyset$
3: **for all** $\langle X, A \rangle \in \Sigma$ **do**
4: $Z \leftarrow \text{AFolding}(\Sigma, A)$
5: $\Sigma^{\text{q}} \leftarrow \Sigma^{\text{q}} \cup \langle Z, A \rangle$
6: **return** Σ^{q}

Algorithm 6 Folding of an attribute w.r.t. an fd set.

```

1: procedure AFOLDING( $\Sigma$ : fd set,  $A$ : attribute)
Require:  $\Sigma$  is parsimonious
Ensure: Returns  $A^{\triangleright}$ , the attribute folding of  $A$  (w.r.t.  $\Sigma$ )
2:    $\Delta \leftarrow \emptyset$   $\triangleright$  consumed fd's
3:    $\Lambda \leftarrow \emptyset$   $\triangleright$  consumed attrs.
4:    $A^* \leftarrow A$   $\triangleright$  stores attrs.  $A$  is found to be 'causally dependent' on (cf. §4.3)
5:    $size \leftarrow 0$ 
6:   while  $size < |A^*|$  do  $\triangleright$  halts when  $A^{(i+1)} = A^{(i)}$ 
7:      $size \leftarrow |A^*|$ 
8:      $\Sigma \leftarrow \Sigma \setminus \Delta$ 
9:     for all  $\langle Y, B \rangle \in \Sigma$  do
10:      if  $B \in A^*$  then
11:         $\Delta \leftarrow \Delta \cup \{\langle Y, B \rangle\}$   $\triangleright$  consumes fd
12:         $A^* \leftarrow A^* \cup Y$ 
13:         $\Lambda \leftarrow \Lambda \cup B$   $\triangleright$  consumes attr.
14:        for all  $C \in Y$  do
15:          if  $C \in \Lambda$  and  $B \in X$  for  $\langle X, C \rangle \in \Delta$  then  $\triangleright$  cyclic fd
16:             $\Lambda \leftarrow \Lambda \setminus B$   $\triangleright$  reingests it to simulate cyclic app. of R5
17:   return  $A^* \setminus \Lambda$ 

```

Theorem 3 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, and Σ an fd set encoded given \mathcal{S} . Now, let $\langle X, A \rangle \in \Sigma$. Then **AFolding**(Σ, A) correctly computes A^{\triangleright} , the attribute folding of A (w.r.t. Σ) in time $O(|\mathcal{S}|^2)$.

Proof 8 For the proof roadmap, note that **AFolding** is monotone (size of A^* can only increase) and terminates precisely when $A^{(i+1)} = A^{(i)}$, where $A^{(i)}$ denotes the attributes in A^* at step i of the outer loop. The folding A^{\triangleright} of A at step i is $A^{(i)} \setminus \Lambda^{(i)}$. We shall prove by induction, given attribute A from fd $X \rightarrow A$ in parsimonious Σ , that $A^* \setminus \Lambda$ returned by **AFolding**(Σ, A) is the unique attribute folding A^{\triangleright} of A . See Appendix, §A.2.2. \square

Remark 4 Let Σ be an arbitrary fd set on attribute set U . Beeri and Bernstein gave a straightforward optimization to (Alg. 4) **XClosure** to make it linear in $|\Sigma| \cdot |U|$ (cf. [63, p. 43-5]), where $|\Sigma| \cdot |U|$ is the maximum length for a string encoding all the fd's. Note that the actual length of such string in our case is exactly $|\mathcal{S}|$. The

optimization mentioned applies likewise to (Alg. 6) **AFolding**.² That is, **AFolding** can be implemented to run in linear time in $|\mathcal{S}|$. \square

Corollary 2 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, and Σ an fd set encoded given \mathcal{S} . Then algorithm **folding**(Σ) correctly computes Σ^{f} , the folding of Σ in time that is $f(|\mathcal{S}|) \Theta(|\mathcal{E}|)$, where $f(|\mathcal{S}|)$ is the time complexity of (Alg. 6) **AFolding**.

Proof 9 See Appendix, §A.2.3. \square

Finally, it shall be convenient to come with a notion of *parsimonious* fd sets (see Def. 13), which is suggestive of a distinguishing feature of such mathematical information systems in comparison with arbitrary information systems.

Def. 13 Let Σ be set of fd's on attributes U . Then, we say that Σ is **parsimonious** if it is canonical and, for all fd's $\langle X, A \rangle \in \Sigma$ with $XA \subseteq U$, there is no $Y \subset U$ such that $Y \neq X$ and $\langle Y, A \rangle \in \Sigma$.

Proposition 3 then shall be useful further in connection with the concept of the folding.

Proposition 3 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, φ a total causal mapping over \mathcal{S} and Σ an fd set encoded through φ given \mathcal{S} . Let Σ^{f} be the folding of Σ , then Σ^{f} is parsimonious.

Proof 10 See Appendix, §A.2.4. \square

4.3. Equivalence with Causal Ordering

Now we show the equivalence of acyclic pseudo-transitive reasoning with causal ordering processing. We start with Theorem 4, which establishes the equivalence between the notion of causal dependency and the fd encoding scheme presented in Chapter 3.

² We omit its tedious exposure here. In short, it shall require one more auxiliary data structure to keep track, for each fd not yet consumed, of how many attributes not yet consumed appear in its rhs.

Theorem 4 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, φ a total causal mapping over \mathcal{S} and Σ an fd set encoded through φ given \mathcal{S} . Then $x_a, x_b \in \mathcal{V}$ are such that x_b is causally dependent on x_a , i.e., $(x_a, x_b) \in C_\varphi^+$ iff there is some non-trivial fd $\langle X, B \rangle \in \Sigma^\triangleright$ with $A \in X$, where $B \mapsto x_b$ and $A \mapsto x_a$.

Proof 11 We prove the statement by induction. We consider first the ‘if’ direction, and then its ‘only if’ converse. See Appendix §A.2.5. \square

Def. 14 then gives useful terminology for a neat concept towards our goal in this chapter.

Def. 14 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a structure with variables $x_a, x_b \in \mathcal{V}$, and φ a total causal mapping over \mathcal{S} inducing set of direct causal dependencies C_φ and its transitive closure C_φ^+ . We say that x_a is a **first cause** of x_b in \mathcal{S} if $(x_a, x_b) \in C_\varphi^+$ and, for no $x \in \mathcal{V}$, we have $(x, x_a) \in C_\varphi^+$.

Proposition 4 connects the notion of first cause with those of exogenous and endogenous variables introduced in Chapter 3.

Proposition 4 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a structure with variable $x \in \mathcal{V}$. Then x can only be a first cause of some $y \in \mathcal{V}$ if x is exogenous. Accordingly, any variable $y \in \mathcal{V}$ can only have some first cause $x \in \mathcal{V}$ if it is endogenous.

Proof 12 Straightforward from definitions, see Appendix §A.2.6. \square

Note that exogenous variables are encoded into fd’s $X \rightarrow A$ with $\phi \in X$. Since the values of such variables are assigned “outside” the system (cf. Remark 3), they are devoid of indirect causal dependencies and then have no uncertainty except for their own. Thus, we have not to be concerned at all with processing the causal (uncertainty) chaining towards them. Our goal is rather find the first causes of the endogenous variables (predictive attributes).

We shall need then the terminology of Def. 15, and then we introduce Lemma 3 paving the way to our goal.

Def. 15 Let \mathcal{S} be a structure, and Σ be a set of fd’s encoded over it. Then $\Upsilon(\Sigma)$, the **v -projection of Σ** , is the subset of fd’s $X \rightarrow A$ such that $v \in X$. Accordingly, $\Phi(\Sigma)$, the **ϕ -projection of Σ** , is the subset of fd’s $X \rightarrow A$ such that $v \notin X$.

$$\begin{array}{lcl}
\Sigma = \{ & \phi & \rightarrow x_1, \\
& \phi & \rightarrow x_2, \\
& \phi & \rightarrow x_3, \\
x_1 x_2 x_3 x_5 v & \rightarrow & x_4, \\
x_1 x_3 x_4 v & \rightarrow & x_5, \\
x_4 v & \rightarrow & x_6, \\
x_5 v & \rightarrow & x_7 \}.
\end{array}
\quad
\Upsilon(\Sigma)^{\text{qv}} = \{
\begin{array}{l}
x_1 x_2 x_3 v x_5 \rightarrow x_4, \\
x_1 x_2 x_3 v x_4 \rightarrow x_5, \\
x_1 x_2 x_3 v x_5 \rightarrow x_6, \\
x_1 x_2 x_3 v x_4 \rightarrow x_7
\end{array}
\}.$$

Figure 4.2. Fd set Σ encoding the structure of Fig. 3.2a and the folding $\Upsilon(\Sigma)^{\text{qv}}$ of its v -projection.

Fig. 4.2 illustrates the v -projection of an fd set and the folding applied over such fd subset in order to compute the first causes of endogenous variables.

Lemma 3 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, φ a total causal mapping over \mathcal{S} and Σ an fd set encoded through φ given \mathcal{S} . Then a variable $x_a \in \mathcal{V}$ can only be a first cause of some variable $x_b \in \mathcal{V}$, where $\langle X, B \rangle \in \Sigma$, and $B \mapsto x_b$, $A \mapsto x_a$, if either (i) $A \in X$ or (ii) $A \notin X$ but there is $\langle Z, C \rangle \in \Sigma^{\triangleright}$ with $A \in Z$ and $C \in X$.

Proof 13 We prove the statement by construction out of Theorem 4, see Appendix §A.2.7. \square

Finally, Theorem 5 further clarifies the purpose of the folding and its meaning in terms of causal ordering.

Theorem 5 Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, φ a total causal mapping over \mathcal{S} and Σ an fd set encoded through φ given \mathcal{S} . Now, let B be an attribute that encodes some variable $x_b \in \mathcal{V}$. If $\langle X, B \rangle \in \Upsilon(\Sigma)^{\text{qv}}$,³ then every first cause x_a of x_b (if any) is encoded by some attribute $A \in X$.

Proof 14 We show that the existence of a missing first cause x_c of x_b for folded $X \xrightarrow{\text{qv}} B$, where $B \mapsto x_b$ and $C \mapsto x_c$ but $C \notin X$ leads to a contradiction. See Appendix §A.2.8. \square

³ Note that the folding is taken w.r.t. the v -projection of Σ , then x_b where $B \mapsto x_b$ is an endogenous variable.

Remark 5 Observe that, on the one hand, the goal of computing the transitive closure C_φ^+ of a set of induced causal dependencies C_φ is to derive the entire causal ordering of a given structure. The goal of folding, on the other hand, is not to discover all variables (attributes) a given variable (attribute) is causally dependent on, but only all of its first causes (if any). \square

In particular, the results just shown comprise a method to compute, for each endogenous variable (predictive attribute), all of its first causes. This is a core goal of the reasoning device developed in this chapter in order to enable the automatic synthesis of hypotheses as uncertain and probabilistic (U-relational) data.

4.4. Experiments

Fig. 4.3 shows the results of experiments we have carried out in order to study how effective the causal reasoning over fd's is in practice, in particular its behavior for hypotheses whose structure \mathcal{S} has been randomly generated over orders of magnitude $|\mathcal{S}| \approx 2^k$, to have length up to $|\mathcal{S}| \approx 2^{20} \lesssim 1M$. The largest structure considered, with $|\mathcal{S}| \approx 1M$, has been generated to have exactly $|\mathcal{E}| = 2.5K$. Like in the experiments of the previous chapter, we have executed ten runs for each tested order of magnitude, and then taken its mean running time in *ms*.⁴

The plot is shown in Fig. 3.8 is in logscale base 2. Notice the linear rate of growth across orders of magnitude (base 2) from $1K$ to $1M$ sized structures. For a growth factor of 2 in structure length (doubled), the time required by causal reasoning grows a factor of 2 (doubled as well). These scalability results are compatible with the computational complexity of **folding**, which is bounded by $O(|\mathcal{S}|^2)$. Yet, that is a bit overestimated time bound as we see in the plot of Fig. 4.3.

4.5. Related Work

The concept of fd set folding and the design of (Alg. 6) **AFolding** as a not quite obvious variant of **XClosure**, is an original approach to the problem of processing the causal ordering of a hypothesis via *acyclic pseudo-transitive reasoning over fd's*. To the best of our knowledge, such a specific form of fd reasoning was an

⁴ The experiments were performed on a 2.3GHZ/4GB Intel Core i5 running Mac OS X 10.6.8.

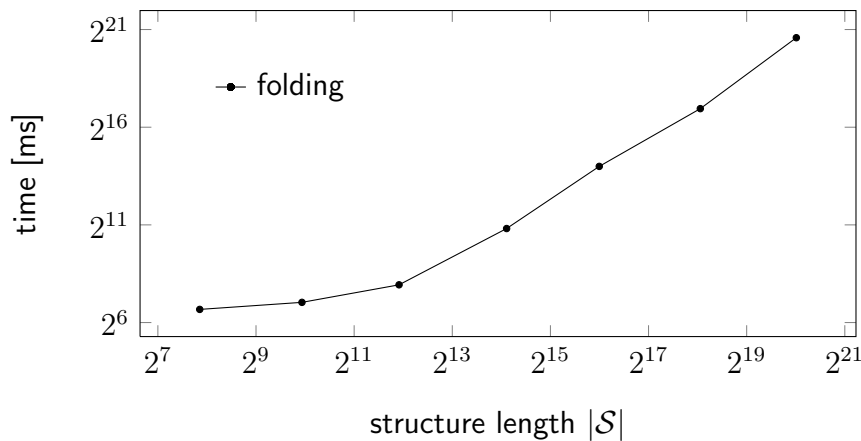


Figure 4.3. Performance of acyclic causal reasoning over fd's (logscale).

yet unexplored problem in the database research literature (reasoning over fd's is extensively covered in Maier [22]).

Recent years have seen the emergence of some foundational work in causality in databases [64]. It is motivated for improving DB usability in terms of providing users with explanations to query answers (and non-answers). Essentially, the idea is borrowed from AI work on causality (cf. §3.6) to identify causal ordering between tuples. Given a query and its result set, the system should be able to explain to the user what tuples 'caused' that answer, or why possibly expected tuples are missing. That requires causal chain of tuples for a given query, which can be computationally expensive as the database instance can be very large. For conjunctive queries, the causality is said to be computed very efficiently [65].

A more specific problem addressed by Kanagal et. al is the so-called sensitivity analysis [66], which is aimed at establishing a more refined connection between the query answer (output) and elements of the DB instance (input) for supporting user interventions. Instead of providing the user with causes and non-causes, the goal is to enable the user to know how changes in the input affect the output. This line of work is strongly related to the vision of 'reverse data management' [67].

Causal reasoning in the presence of constraints (viz., fd's) is an yet unexplored topic, though called for as worth of future work by Meliou et al. [64, p. 3]. The fd's are rich information that can be exploited for the sake of explanation and sensitivity analysis. Once they are available, it is intuitive that the search space of such problems shall be significantly reduced.

In fact, our encoding of equations into fd's captures the causal chain from exogenous (input) to endogenous (output) tuples in at schema level. Nonetheless our form of causal reasoning over fd's is geared for hypothesis management and analytics, from a uncertainty management point of view. A concrete connection to causality in DB's is not yet established.

4.6. Summary of Results

In this chapter we have studied and developed a technique for acyclic causal reasoning over fd's. We list the results achieved as follows.

- We have developed principled concepts and a core algorithm, viz., (Alg. 5) the **folding**, in order to perform acyclic pseudo-transitive reasoning over fd's. This is towards an efficient method for causal reasoning, yet elegant as a database formalism for the systematic construction of hypothesis probabilistic DB's.
- We have given a reasonably tight time bound for the behavior of such reasoning device in terms of the structure given as input. We have established (cf. Theorem 3, Corollary 2) the time bound of $O(|\mathcal{S}|^2)$ for the **folding** algorithm.
- We have shown the correctness of the **folding** algorithm in connection with causal reasoning (cf. Theorem 4, Theorem 5).
- We have defined the core notion of first causes (cf. Def. 14, Proposition 4), which is meant to guide the procedure of **U-intro** (Chapter 5) by precisely capturing the uncertainty factors on endogenous variables (predictive attributes). This is similar, yet markedly different from computing the transitive closure of causal dependencies (cf. Remark 5).
- We have performed experiments (cf. Fig. 4.3) to study how effective the approach of causal reasoning over fd's is in practice, or how it scales for hypotheses whose structure \mathcal{S} is randomly generated to have length up to the order of $|\mathcal{S}| \lesssim 1M$. The experiments show that the time bound, though already effective for very large structures, are a bit overestimated.

Chapter 5

Probabilistic Database Synthesis

In this chapter we present a technique to synthesize hypothesis U-relations. At this stage of the pipeline, relational schema \mathbf{H} is loaded with datasets computed from the hypotheses under alternative trials (input settings). The challenge is how to model or design its probabilistic version (i.e., render the U-relations \mathbf{Y}) so that it is suitable for data-driven hypothesis management and analytics.

In §5.1 we introduce U-relational DB's. In §5.2 we present a running example to illustrate the uncertainty introduction procedure (U-intro in the pipeline, cf. Fig. 1.4). In §5.3 we present the technique to factorize the uncertainty present in the 'big' fact table in terms of the well-defined uncertainty factors. In §5.4 then we show how to propagate such uncertainty into the predictive attributes properly, i.e., based on their first causes detected as shown in Chapter 4. In §5.6, we discuss related work. Finally, in §5.7 we conclude the chapter.¹

5.1. Preliminaries: U-Relations and Probabilistic WSA

Three remarkable features of U-relations are: *expressiveness* (being closed under positive relational algebra queries); *succinctness* (efficient storage of a very large number of possible worlds through vertical decompositions to support attribute-level uncertainty); and *efficient query processing* (including confidence computation) [18].

A U-relational database or U-database is a finite set of structures,

$$\mathbf{W} = \{ \langle R_1^1, \dots, R_m^1, p^{[1]} \rangle, \dots, \langle R_1^n, \dots, R_m^n, p^{[n]} \rangle \},$$

¹ We postpone the presentation of experiments on p-DB synthesis (U-intro as a whole) to §6.4.

of relations R_1^i, \dots, R_m^i and numbers $0 < p^{[i]} \leq 1$ such that $\sum_{1 \leq i \leq n} p^{[i]} = 1$. An element $R_1^i, \dots, R_m^i, p^{[i]} \in \mathbf{W}$ is a *possible world*, with $p^{[i]}$ being its probability [18].

Probabilistic world-set algebra (p-WSA) consists of the operations of relational algebra, an operation for computing tuple confidence **conf**, and the **repair-key** operation for introducing uncertainty — by giving rise to alternative worlds as maximal-subset repairs of an argument key [18].

Let $R_\ell[U]$ be a relation, and $XA \subseteq U$. For each possible world $\langle R_1, \dots, R_m, p \rangle \in \mathbf{W}$, let $A \in U$ contain only numerical values greater than zero and let R_ℓ satisfy the fd $(U \setminus A) \rightarrow U$. Then, **repair-key** is:

$$\llbracket \text{repair-key}_{X@A}(R_\ell) \rrbracket(\mathbf{W}) := \left\{ \langle R_1, \dots, R_\ell, R_m, \hat{R}_\ell[U \setminus A], \hat{p} \rangle \right\},$$

where $\langle R_1, \dots, R_\ell, R_m, p \rangle \in \mathbf{W}$, \hat{R}_ℓ is a maximal repair of fd $X \rightarrow U$ in R_ℓ , and $\hat{p} = p \cdot \prod_{t \in \hat{R}_\ell} \frac{t.B}{\sum_{s \in R_\ell : s.X = t.X} s.B}$.

U-relations (cf. Fig. 5.1) have in their schema a set of pairs (V_i, D_i) of *condition columns* (cf. [18]) to map each discrete random variable x_i to one of its possible values (e.g., $x_1 \mapsto 1$). The world table W stores their marginal probabilities (cf. the notion of *pc-tables* [17, Ch. 2]). For an illustration of the data transformation from certain to uncertain relations, consider query (5.1) in p-WSA's extension of relational algebra, whose result set is materialized into U-relation Y_0 as shown in (Fig. 5.1).

$$Y_0 := \pi_{\phi, v}(\text{repair-key}_{\phi@Conf}(H_0)). \quad (5.1)$$

Also, let $R[\overline{V_i D_i} \mid sch(R)]$, $S[\overline{V_j D_j} \mid sch(S)]$ be two U-relations, where $R.\overline{V_i D_i}$ is the union of all pairs of condition columns $V_i D_i$ in R , then operations of selection $\llbracket \sigma_\psi(R) \rrbracket$, projection $\llbracket \pi_Z(R) \rrbracket$, and product $\llbracket R \times S \rrbracket$ issued in relational algebra are rewritten in positive relational algebra on U-relations:

$$\llbracket \sigma_\psi(R) \rrbracket := \sigma_\psi(R[\overline{V_i D_i} \mid sch(R)]);$$

$$\llbracket \pi_Z(R) \rrbracket := \pi_{\overline{V_i D_i} Z}(R);$$

$$\llbracket R \times S \rrbracket := \pi_{(R.\overline{V_i D_i} \cup S.\overline{V_j D_j}) \rightarrow \overline{V D} \cup sch(R) \cup sch(S)}(R \bowtie_{R.\overline{V_i D_i} \text{ is consistent with } S.\overline{V_j D_j}} S).$$

H_0	ϕ	v	Conf
	1	1	2
	1	2	2
	1	3	1

Y_0	$V \mapsto D$	ϕ	v	W	$V \mapsto D$	Pr
	$x_0 \mapsto 1$	1	1		$x_0 \mapsto 1$.4
	$x_0 \mapsto 2$	1	2		$x_0 \mapsto 2$.4
	$x_0 \mapsto 3$	1	3		$x_0 \mapsto 3$.2

Figure 5.1. U-relation generated by the repair-key operation.

If R and S have k and ℓ pairs of condition columns each, then $\llbracket R \times S \rrbracket$ returns a U-relation with $k + \ell$ such pairs. If $k = 0$ or $\ell = 0$ (or both), then R or S (or both) are classical relations, but the rewrite rules above apply accordingly. All that rewriting is parsimonious translation (sic. [18]): the number of algebraic operations does not increase and each of the operations selection, projection and product/join remains of the same kind. Query plans are hardly more complicated than the input queries. In fact, it has been verified that off-the-shelf relational database query optimizers do well in practice.

For a comprehensive overview of U-relations and p-WSA we refer the reader to [18]. In this thesis we look at U-relations from the point of view of p-DB design, for which no methodology has yet been proposed. We are concerned in particular with hypothesis management applications [10].

5.2. Running Example

Before proceeding, we consider Example 6, which is fairly representative to illustrate how to deal with correlations in the predictive data of deterministic hypotheses for the sake of suitable data-driven analytics.

Example 6 We explore three slightly different theoretical models in population dynamics with applications in Ecology, Epidemics, Economics, etc: (5.2) Malthus' model, (5.3) the logistic equation and (5.4) the Lotka-Volterra model. In practice, such equations are meant to be extracted from MathML-compliant XML files (cf. Chapter 6). For now, consider that the ordinary differential equation notation ' \dot{x} '

is read ‘variable x is a function of time t given initial condition x_0 .’

$$\dot{x} = bx \tag{5.2}$$

$$\dot{x} = b(1 - x/K)x \tag{5.3}$$

$$\begin{cases} \dot{x} &= x(b - py) \\ \dot{y} &= y(rx - d) \end{cases} \tag{5.4}$$

The models are completed (by the user) with additional equations to provide the values of exogenous variables (or ‘input parameters’),² e.g., $x_0 = 200$, $b = 10$, such that we have SEM’s (resp.) $\mathcal{S}_k(\mathcal{E}_k, \mathcal{V}_k)$ for $k = 1..3$,

- $\mathcal{E}_1 = \{ f_1(t), f_2(x_0), f_3(b), f_4(x, t, x_0, b) \}$;
- $\mathcal{E}_2 = \{ f_1(t), f_2(x_0), f_3(K), f_4(b), f_5(x, t, x_0, K, b) \}$;
- $\mathcal{E}_3 = \{ f_1(t), f_2(x_0), f_3(b), f_4(p), f_5(y_0), f_6(d), f_7(r), f_8(x, t, x_0, b, p, y), f_9(y, t, y_0, d, r, x) \}$.

Fig. 5.2 shows the fd sets encoded from structures \mathcal{S}_k given above.³ We also consider trial datasets for hypothesis $v=3$ (viz., the Lotka-Volterra model), which are loaded into the ‘big’ fact table relation $H_3 \in \mathbf{H}$ as shown in Fig. 5.3. We admit a special attribute ‘trial id’ tid to keep hypothesis trials identified until uncertainty is introduced in a controlled way by p-DB synthesis (U-intro stage, cf. Fig. 1.4).

□

² Given $\mathcal{S}(\mathcal{E}, \mathcal{V})$, it is actually a task of the encoding algorithm (viz., sub-procedure TCM) to infer, for each variable $x \in \mathcal{V}$, whether it is exogenous or endogenous by means of the total causal mapping.

³ Recall that domain variables like time t are informed to the encoding algorithm to suppress an fd $\phi \rightarrow t$.

$\Sigma_1 = \{ \phi \rightarrow x_0,$ $\phi \rightarrow b,$ $x_0 b t v \rightarrow x \}.$	$\Sigma_3 = \{ \phi \rightarrow x_0,$ $\phi \rightarrow b,$ $\phi \rightarrow p,$ $\phi \rightarrow y_0,$ $\phi \rightarrow d,$ $\phi \rightarrow r,$ $x_0 b p t v y \rightarrow x,$ $y_0 d r t v x \rightarrow y \}.$
$\Sigma_2 = \{ \phi \rightarrow x_0,$ $\phi \rightarrow K,$ $\phi \rightarrow b,$ $x_0 K b t v \rightarrow x \}.$	

Figure 5.2. Fd sets encoded from the given structures $\mathcal{S}_k(\mathcal{E}_k, \mathcal{V}_k)$ for hypotheses $k = 1..3$ from Example 6.

H ₃	tid	ϕ	v	t	x_0	b	p	y_0	d	r	x	y
	1	1	3	0	30	.5	.02	4	.75	.02	30	4
	1	1	3	...	30	.5	.02	4	.75	.02
	2	1	3	0	30	.5	.018	4	.75	.023	30	4
	2	1	3	...	30	.5	.018	4	.75	.023
	3	1	3	0	30	.4	.02	4	.8	.02	30	4
	3	1	3	...	30	.4	.02	4	.8	.02
	4	1	3	0	30	.4	.018	4	.8	.023	30	4
	4	1	3	...	30	.4	.018	4	.8	.023
	5	1	3	0	30	.397	.02	4	.786	.02	30	4
	5	1	3	...	30	.397	.02	4	.786	.02
	6	1	3	0	30	.397	.018	4	.786	.023	30	4
	6	1	3	5	30	.397	.018	4	.786	.023	50.1	62.9
	6	1	3	10	30	.397	.018	4	.786	.023	13.8	8.65
	6	1	3	15	30	.397	.018	4	.786	.023	79.3	8.23
	6	1	3	20	30	.397	.018	4	.786	.023	12.6	30.7
	6	1	3	...	30	.397	.018	4	.786	.023

Figure 5.3. ‘Big’ fact table H_3 of hypothesis $k=3$ from Example 6 loaded with trial datasets identified by special attribute tid.

Given the ‘big’ fact table H_3 , p-DB synthesis has two main parts: process the ‘empirical’ uncertainty present in the ‘big’ fact table and synthesize it out (decompose it) into independent u-factors (*u-factorization*); and then propagate it precisely into the predictive data (*u-propagation*).

5.3. U-Factorization

As we have seen in §5.1, the **repair-key** operation allows one to create a discrete random variable in order to repair an argument key in a given relation. Our goal here is to devise a technique to perform such operation in a principled way for hypothesis management. It is a basic design principle to have exactly one random variable for each distinct uncertainty factor ('u-factor' for short), which requires carefully identifying the actual sources of uncertainty present in relations \mathbf{H} .

The multiplicity of (competing) hypotheses is itself a standard one, viz., the *theoretical* u-factor. Consider an 'explanation' table like H_0 in Fig. 5.1, which stores (as foreign keys) all hypotheses available and their target phenomena. We can take such H_0 as explanation table for the three hypotheses of Example 6. Then a discrete random variable V_0 is defined into $Y_0 [V_0 D_0 | \phi v]$ by query formula (5.1). U-relation Y_0 is considered standard in p-DB synthesis, as the repair of ϕ as a key in (standard) H_0 .

Hypotheses, nonetheless, are (abstract) 'universal statements' [15]. In order to produce a (concrete) valuation over their endogenous attributes (predictions), one has to inquire into some particular 'situated' phenomenon ϕ and tentatively assign a valuation over the exogenous attributes, which can be eventually tuned for a target ϕ . The multiplicity of such (competing) empirical estimations for a hypothesis k leads to Problem 1, viz., learning *empirical* u-factors for each $H_k \in \mathbf{H}$.

Problem 1 Let Σ_k be an fd set encoded given hypothesis structure \mathcal{S}_k , and H_k its 'big' fact table relation loaded with trial data. Now, let Z be the set of attributes encoding exogenous variables in H_k , then the problem of **u-factor learning** is:

- (1) to **infer** in H_k 'casual' fd's $\phi B_i \rightarrow B_j$, $\phi B_j \rightarrow B_i$ not in Σ_k (strong input correlations), where $B_i, B_j \in Z$;
- (2) to form maximal **groups** $G_1, \dots, G_n \subseteq Z$ of attributes such that for all $B_i, B_j \in G_a$, the casual fd's $\phi B_i \rightarrow B_j$ and $\phi B_j \rightarrow B_i$ hold in H_k ;
- (3) to pick, for each group G_a , any $A \in G_a$ as a **pivot** representative and insert $\phi A \rightarrow B$ into an fd set Ω_k for all $B \in (G_a \setminus A)$.

H_3	tid	ϕ	v	t	x_0	b	p	y_0	d	r	x	y
	1	1	3	0	30	.5	.02	4	.75	.02	30	4
	1	1	3	...	30	.5	.02	4	.75	.02
	2	1	3	0	30	.5	.018	4	.75	.023	30	4
	2	1	3	...	30	.5	.018	4	.75	.023
	3	1	3	0	30	.4	.02	4	.8	.02	30	4
	3	1	3	...	30	.4	.02	4	.8	.02
	4	1	3	0	30	.4	.018	4	.8	.023	30	4
	4	1	3	...	30	.4	.018	4	.8	.023
	5	1	3	0	30	.397	.02	4	.786	.02	30	4
	5	1	3	...	30	.397	.02	4	.786	.02
	6	1	3	0	30	.397	.018	4	.786	.023	30	4
	6	1	3	5	30	.397	.018	4	.786	.023	50.1	62.9
	6	1	3	10	30	.397	.018	4	.786	.023	13.8	8.65
	6	1	3	15	30	.397	.018	4	.786	.023	79.3	8.23
	6	1	3	20	30	.397	.018	4	.786	.023	12.6	30.7
	6	1	3	...	30	.397	.018	4	.786	.023

Figure 5.4. ‘Big’ fact table H_3 of hypothesis $k=3$ from Example 6 with u-factors $\{b, d\}$ and $\{p, r\}$ emphasized (resp.) in colors green and red.

U-factor learning is meant to process only the attributes $Z \subset U$ from $H_k[U]$ that are inferred exogenous in the given hypothesis, i.e., for all $A \in Z$, there is an fd $\langle X, A \rangle \in \Phi(\Sigma_k)$, where the latter is the ϕ -projection of Σ_k . Such attributes are then ‘officially’ unrelated. In fact, by ‘casual’ fd’s we mean correlations that, for a set of experimental trials, may occasionally show up in the trial input data; e.g., $x_0 \leftrightarrow y_0$ hold in H_3 , but not because x_0 and y_0 are related in principle (theory).

Fig. 5.4 helps to illustrate Problem 1 through the ‘big’ fact table. We emphasize u-factors $\{b, d\}$ and $\{p, r\}$ in colors green and red. Observe that values of b are strongly correlated (one-to-one) with values of d for $\phi = 1$, just like p and r . Note also that $\{x_0, y_0\}$ can be seen as a certain factor. From the user point of view, this is a record that reflects a common practice in computational science known as (parameter) sensibility analysis.

Problem 1 is dominated by the (problem of) discovery of fd’s in a relation, which is not really a new problem (e.g., see [68]). We then keep focus on the synthesis method as a whole and omit our detailed u-factor-learning algorithm in

$\Omega_3 = \{ \begin{array}{l} \phi x_0 \rightarrow y_0, \\ \phi b \rightarrow d, \\ \phi p \rightarrow r, \\ x_0 b p t v y \rightarrow x, \\ y_0 d r t v x \rightarrow y \end{array} \}.$	$\Omega_3^{\nabla} = \{ \begin{array}{l} \phi x_0 \rightarrow y_0, \\ \phi b \rightarrow d, \\ \phi p \rightarrow r, \\ \phi x_0 b p t v y \rightarrow x, \\ \phi x_0 b p t v x \rightarrow y \end{array} \}.$
---	--

Figure 5.5. Fd set Ω_3 (compare with Σ_3) and its folding Ω_3^{∇} .

particular.⁴ Its output, fd set Ω_k , is then filled in (completed) with the v -projection $\Upsilon(\Sigma_k)$. For illustration consider hypothesis $v = 3$ and its trial input data recorded in H_3 in Fig. 5.3. We show its resulting fd set Ω_3 in Fig. 5.5 (left), together with its folding Ω_k^{∇} (right). The latter is then input to (Alg. 7) **merge** to get the final information necessary for the actual synthesis of U-relations, as captured in Def. 16. For an illustration of the merging of fd's with equivalent left-hand sides, note in Fig. 5.5 (right) that $\phi x_0 b p t v y \leftrightarrow \phi x_0 b p t v x$ holds in $(\Omega_3^{\nabla})^+$.

Def. 16 Let \mathcal{S}_k and H_k be the complete structure and ‘big’ fact table of hypothesis k , and Σ_k an fd set defined $\Sigma_k \triangleq \text{h-encode}(\mathcal{S}_k)$. Now, let $\Omega_k \triangleq \text{u-factor-learning}(H_k, \Phi(\Sigma_k)) \cup \Upsilon(\Sigma_k)$, and Ω_k^{∇} be the folding of Ω_k . Finally, define $\Gamma_k \triangleq \text{merge}(\Omega_k^{\nabla})$. Then we say that Γ_k is the **u-factorization** of \mathcal{S}_k over H_k .

Algorithm 7 Merge fd's with equivalent left-hand sides.

```

1: procedure MERGE( $\Sigma$  : fd set)
2:    $\Omega \leftarrow \emptyset$ 
3:   for all  $\langle X, C \rangle \in \Sigma$  do
4:     if there is  $\langle Z, W \rangle \in \Omega$  such that  $X \leftrightarrow Z$  holds in  $\Sigma^+$  then
5:        $\Omega \leftarrow \Omega \setminus \langle Z, W \rangle$ 
6:        $S \leftarrow X \setminus Z$ 
7:        $\Omega \leftarrow \Omega \cup \langle Z, W S C \rangle$  ▷ merges equivalent keys
8:     else
9:        $\Omega \leftarrow \Omega \cup \langle X, C \rangle$ 
10:  return  $\Omega$ 

```

⁴ In short, we make use of relational algebra **group-by** operation and build a pruned lattice of attribute groups having the same number of rows under the grouping (similarly to [68]).

Remark 6 Let Γ_k be the u-factorization of structure \mathcal{S}_k over ‘big’ fact table H_k . Then every fd in Γ_k encodes a clear-cut *claim*, either empirical, in $\Phi(\Gamma_k)$, or theoretical, in $\Upsilon(\Gamma_k)$. That is ensured by the **merge** algorithm, which groups fd’s in Ω_k^{\rightarrow} with equivalent left-hand sides. \square

We are now able to employ a notion of u-factor decomposition formulated in Def. 17 into query formula (5.5) in p-WSA’s extension of relational algebra.

Def. 17 Let \mathcal{S}_k be the complete structure of hypothesis k , and $H_k[U]$ its ‘big’ fact table such that Γ_k is the u-factorization of \mathcal{S}_k over H_k . Now, let $G_a \subset U$ be a set of attributes $G_a = A G$ such that, for all $B \in G$, an fd $\phi A \rightarrow B$ exists in $\Phi(\Gamma_k)$. Then we define U-relation $Y_k^i[V_i D_i | \phi A G]$ by query formula (5.5), and say that Y_k^i is a **u-factor projection** of H_k .

$$Y_k^i := \pi_{\phi A G}(\text{repair-key}_{\phi @ \text{count}}(\gamma_{\phi, A, G, \text{count}(*)}(H_k))) \quad (5.5)$$

where γ is relational algebra’s grouping operator.

Y_3^1	$V \mapsto D$	ϕ	x_0	y_0	Y_3^2	$V \mapsto D$	ϕ	b	d	Y_3^3	$V \mapsto D$	ϕ	p	r
	$x_2 \mapsto 1$	2	30	4		$x_3 \mapsto 1$	2	.5	.5		$x_4 \mapsto 1$	2	.020	.020
						$x_3 \mapsto 2$	2	.4	.8		$x_4 \mapsto 2$	2	.018	.023
						$x_3 \mapsto 3$	2	.397	.786					

Figure 5.6. U-factor projections rendered for hypothesis $v = 3$.

The synthesis of u-factor projections, in particular the application of **repair-key** (cf. Eq. 5.5), has an important consequence for the u-factorization Γ_k of H_k , viz., the introduction of new fd’s into Γ'_k defined as follows (see Def. 18). We shall consider it (rather than Γ_k) to study design-theoretic properties of synthesized \mathbf{Y}_k in §5.5.

Def. 18 Let \mathcal{S}_k and $H_k[U]$ be (resp.) the complete structure and ‘big’ fact table of hypothesis k , and Γ_k be the u-factorization of \mathcal{S}_k over H_k . Now, let $\Gamma'_k \triangleq \bigcup_{i \in I} \{\phi \rightarrow A_i G_i\} \cup \Gamma_k$, where I indexes all u-factor projections $Y_k^i[V_i D_i | \phi A_i G_i]$ of H_k . We say that Γ'_k is the **repaired factorization** of \mathcal{S}_k over H_k .

5.4. U-Propagation

U-propagation is a central part of U-intro and the pipeline itself. Recall that all the machinery developed so far, from hypothesis encoding to causal reasoning to u-factorization is for enabling predictive analytics. Let us briefly reconstruct it.

For hypothesis structure $\mathcal{S}_k(\mathcal{E}, \mathcal{V})$, take any endogenous variable $x_c \in \mathcal{V}$ encoded by attribute $C \mapsto x_c$. There should be exactly one fd $\langle X, C \rangle \in \Upsilon(\Sigma_k)^{\text{q}}$. By Theorem 5, for every first cause x_b of x_c there is $B \in X$ where $B \mapsto x_b$ with $x_b \in \mathcal{V}$. Now, observe that when Ω_k is rendered by u-factor learning, it is filled partly with fd's from $\Upsilon(\Sigma_k)$, and partly with fd's processed from $\Phi(\Sigma_k)$. This is to summarize exogenous variables into clear-cut independent u-factors. It means that, after u-factor learning, each first cause x_b encoded by B shall be represented by some pivot attribute A_i which is, if not $A_i = B$ itself, then occasionally strongly correlated to it (i.e., $\phi A_i \rightarrow B$ and $\phi B \rightarrow A_i$ hold in H_k).

Further then Ω_k is subject to folding such that, if $B \in X$ for $\langle X, C \rangle \in \Upsilon(\Omega_k)$, now we have $A_i \in Z$ for $\langle Z, C \rangle \in \Upsilon(\Omega_k^{\text{q}})$. This processing from fd $X \rightarrow C$ (where X contains the first causes) into $Z \rightarrow C$ (where Z contains only their pivot representatives) is meant for enabling an economical representation of uncertainty. Our running example is small, but such a principle is quite relevant for large-scale hypotheses (say, when $|\mathcal{S}| \approx 1M$). The correctness of such u-factor summarization shall be ensured by Proposition 3, which let us know that Ω_k^{q} is parsimonious then (by Def. 13) canonical, therefore (by Def. 9) left-reduced.

Now, all fd's in $\Upsilon(\Omega_k^{\text{q}})$ have form $Z \rightarrow C$ and we are almost ready for u-propagation. Note that, as a result of u-factorization, each pivot attribute $A_i \in Z$ is associated with random variable V_i from U-relation $Y_k^i [V_i D_i | \phi A_i G_i]$. Then we shall use each $A_i \in Z$ (from $Z \rightarrow C$) as a surrogate to Y_k^i in order to propagate factorized uncertainty into 'predictive' U-relations $Y_k^j [\overline{V_j D_j} | S T]$ by a join formula. Attribute sets S and T are defined after merging fd's in Ω_k^{q} with equivalent lhs to get Γ_k and pass it as argument for synthesis. We let $S = Z \setminus W$ such that S contains the domain variables only (e.g., ϕ, v, t). The pivot attributes in Z shall not be included in the data columns of Y_k^j , but leave their trace through the condition columns $\overline{V_j D_j}$ that annotate $\text{sch}(Y_k^j)$ as a repair of the key $S \rightarrow T$.

All that (cf. Def. 19) is abstracted into general p-WSA query formula (5.6), and employed in (Alg. 8) **synthesize** to accomplish u-propagation (Part II).

Def. 19 Let \mathcal{S}_k be the complete structure of hypothesis k , and $H_k[U]$ its ‘big’ fact table such that Γ_k is the u-factorization of \mathcal{S}_k over H_k . Now, let $Z \rightarrow T$ be an fd in $\Upsilon(\Gamma_k)$. Then we define U-relation $Y_k^j [\overline{V_j D_j} \mid S T]$ by query formula (5.6), and say that Y_k^j is a **predictive projection** of H_k , where:

$$Y_k^j := \pi_{S,T}(\sigma_{v=k}(Y_0) \bowtie (\bowtie_{i \in I} Y_k^i) \bowtie H_k) \quad (5.6)$$

- (a) $Y_k^i [V_i D_i \mid \phi A_i G_i]$ is a u-factor projection of H_k ;
- (a) we have $i \in I$ if $A_i \in Z$;
- (c) we take $S = Z \setminus W$, where W is the set of all pivot attributes representing first causes in Z .

Algorithm 8 p-DB synthesis applied over folding fd set.

1: **procedure** SYNTHESIZE(\mathcal{S}_k : structure, H_k : ‘big’ table, Y_0 : explan. table)

Require: \mathcal{S}_k is complete

Ensure: U-relations \mathbf{Y}_k returned are a BCNF, lossless decomposition of H_k

- 2: $\Sigma_k \leftarrow \text{h-encode}(\mathcal{S}_k)$
 - 3: $\Omega_k \leftarrow \text{u-factor-learning}(\Phi(\Sigma_k), H_k) \cup \Upsilon(\Sigma_k)$
 - 4: $\Gamma_k \leftarrow \text{merge}(\text{folding}(\Omega_k))$
-

Part I: **U-factorization**

- 5: **for all** $\langle \phi A, G \rangle \in \Phi(\Gamma_k)$ **do** ▷ scans over the u-factors of hypothesis k
 - 6: $Y_k^i \leftarrow \pi_{\phi, A, G}(\text{repair-key}_{\phi @ \text{count}}(\gamma_{\phi, A, G, \text{count}(*)}(H_k)))$
 - 7: $\mathbf{Y}_k \leftarrow \mathbf{Y}_k \cup Y_k^i$
-

Part II: **U-propagation**

- 8: **for all** $\langle Z, T \rangle \in \Upsilon(\Gamma_k)$ **do** ▷ scans over the claims of hypothesis k
 - 9: $W \leftarrow \emptyset$ ▷ prepares to keep track of u-factor pivot attributes
 - 10: **for all** $Y_k^i [\phi A_i G_i] \in \mathbf{Y}_k$ **do**
 - 11: **if** $A \in Z$ **then** ▷ A is a first cause of all $B \in T$
 - 12: $I = I \cup \{i\}$ ▷ indexes the u-factor projection
 - 13: $W \leftarrow W \cup A$ ▷ keeps track of u-factor’s pivot attribute
 - 14: $S \leftarrow Z \setminus W$ ▷ removes u-factor pivot attributes
 - 15: $Y_k^j \leftarrow \pi_{S,T}(\sigma_{v=k}(Y_0) \bowtie (\bowtie_{i \in I} Y_k^i) \bowtie H_k)$
 - 16: $\mathbf{Y}_k \leftarrow \mathbf{Y}_k \cup Y_k^j$
 - 17: **return** \mathbf{Y}_k
-

Fig. 5.7 shows the rendered U-relations for hypothesis $k = 3$ whose ‘big’ fact table is shown in Fig. 5.3. Note that $\text{tid} = 6$ in H_3 corresponds now to $\theta = \{x_1 \mapsto 3, x_2 \mapsto 1, x_3 \mapsto 3, x_4 \mapsto 2\}$, where θ defines a particular world in \mathbf{W} whose probability is $\Pr(\theta) \approx .055$. This value is derived from the marginal probabilities stored in world table W (see Fig. 5.7) as a result of the application of formulas Eq. 5.1 and Eq. 5.5.

Remark 7 Observe that, although (Alg. 8) `synthesize` operates locally for each hypothesis k , the effects of p-DB synthesis (U-intro) in the pipeline are global on account of the (global) ‘explanation’ relation H_0 (then U-relation Y_0), e.g., see Fig. 5.7. In fact, the probability of each tuple (row), say, in U-relation Y_k^j with $\phi = p$ for hypothesis $v = k$, is distributed among all the hypotheses $\ell \neq k$ that are keyed in Y_0 under $\phi = p$, i.e., all hypotheses that compete at $\phi = p$. \square

				W			$V \mapsto D$		Pr	
							$x_1 \mapsto 1$.4		
							$x_1 \mapsto 2$.4		
							$x_1 \mapsto 3$.2		
							$x_2 \mapsto 1$	1		
							$x_3 \mapsto 1$.33		
							$x_3 \mapsto 2$.33		
							$x_3 \mapsto 3$.33		
							$x_4 \mapsto 1$.5		
							$x_4 \mapsto 2$.5		

Y_0		$V \mapsto D$		ϕ	v
		$x_1 \mapsto 1$	1	1	
		$x_1 \mapsto 2$	1	2	
		$x_1 \mapsto 3$	1	3	

Y_3^1		$V \mapsto D$		ϕ	x_0	y_0
		$x_2 \mapsto 1$	1	30	4	

Y_3^2		$V \mapsto D$		ϕ	b	d
		$x_3 \mapsto 1$	1	.5	.5	
		$x_3 \mapsto 2$	1	.4	.8	
		$x_3 \mapsto 3$	1	.397	.786	

Y_3^3		$V \mapsto D$		ϕ	p	r
		$x_4 \mapsto 1$	1	.020	.020	
		$x_4 \mapsto 2$	1	.018	.023	

Y_3^4		$V_1 \mapsto D_1$	$V_2 \mapsto D_2$	$V_3 \mapsto D_3$	$V_4 \mapsto D_4$	ϕ	v	t	y	x
		$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 1$	$x_4 \mapsto 1$	1	3	1900	4	30
		$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 1$	$x_4 \mapsto 1$	1	3
		1	3
		$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 3$	$x_4 \mapsto 2$	1	3	1900	4	30
		$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 3$	$x_4 \mapsto 2$	1	3	1901	4.12	41.5
		$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 3$	$x_4 \mapsto 2$	1	3	1902	5.78	56.7
		$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 3$	$x_4 \mapsto 2$	1	3	1903	11.7	72.8
		$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 3$	$x_4 \mapsto 2$	1	3	1904	31.1	75.9
		$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 3$	$x_4 \mapsto 2$	1	3

U-relations rendered by p-DB synthesis are ready for querying. Typical queries comprise the $\text{conf}()$ aggregate operation, inquiring the probability (or confidence) for each tuple to ‘true’ in the probability space captured by the hypothesis competition. We illustrate queries in Chapter 6.

5.5. Design-Theoretic Properties

For the U-intro procedure to be meaningful, we have to study design-theoretic properties of u-factor projections and prediction projections synthesized out of ‘big’ fact table H_k for the sake of predictive analytics. In particular, for the projections to be claim-centered, we submit that they should satisfy Boyce-Codd normal form (BCNF, cf. Def. 21) w.r.t. the repaired factorization Γ'_k of H_k ; and for them to be a correct decomposition of the uncertainty present in H_k , their join should be lossless (preserve the data in H_k , cf. Def. 22) w.r.t. Γ'_k .

Note that in this study we consider repaired factorization Γ'_k (not u-factorization Γ_k), since it is the one which actually holds in \mathbf{Y}_k after key repairing.

5.5.1 Claim-Centered Decomposition

As emphasized through Remark 6, every fd in u-factorization Γ_k is a claim (cf. Remark 6), then the same holds for repaired factorization Γ'_k . Thus, for a claim-centered decomposition of ‘big’ fact table H_k , it is desirable that U-relational schema \mathbf{Y}_k that it satisfies BCNF w.r.t. Γ'_k . BCNF (‘do not represent the same fact twice’ [21, p. 251]) is our notion of ‘good design’ for uncertainty decomposition in view of predictive analytics. This is to avoid the uncertainty of one claim to be undesirably mixed with the uncertainty of another claim.

Def. 20 Let $R[U]$ be a relation scheme over set U of attributes, and Σ a set of fd’s. Then the **projection** of Σ onto $R[U]$, written $\pi_U(\Sigma)$, is the subset $\Sigma' \subseteq \Sigma$ of fd’s $X \rightarrow Z$ such that $XZ \subseteq U$.

Def. 21 Let $R[U]$ be a relation scheme over set U of attributes, and Σ a set of fd’s on U . We say that:

- (a) R is in **BCNF** if, for all $\langle X, A \rangle \in \Sigma^+$ with $A \notin X$ and $XA \subseteq U$, we have $X \rightarrow U$ (i.e., X is a superkey for R);

(b) A schema \mathbf{R} is in BCNF if all of its schemes $R_1, \dots, R_n \in \mathbf{R}$ are in BCNF.

Example 7 To illustrate the concept of BCNF, let us consider canonical fd set $\Sigma = \{A \rightarrow B, B \rightarrow C\}$ over attributes $U = \{A, B, C\}$, and a tentative schema containing a single relation $R[ABC]$. This relation is not in BCNF because, for one, $B \rightarrow C$ violates it ($C \not\subseteq B$ but B is not a superkey for R). \square

Observe also that an overdecomposed schema may (trivially) satisfy BCNF. For example, let $\Sigma = \{A \rightarrow B, A \rightarrow C\}$ then by Def. 21 both schemas $\mathbf{R} = R_1[ABC]$ and $\mathbf{R}' = \{R_1[AB], R_2[AC]\}$ are in BCNF w.r.t. Σ . The second, however, breaks data into two tables making their access more difficult than necessary since both B and C brings in information about A . That is, if the schema were to be synthesized over the fd's in Σ , then it would be desirable to apply (R4) union or merge them before. Our point is that, if we target at a BCNF-satisfying schema, then it is also desirable for it to be the minimal-cardinality schema in BCNF.

Theorem 6 guarantees the BCNF property w.r.t. Γ'_k by design for every schema \mathbf{Y}_k rendered by (Alg. 8) `synthesize` over Γ_k .

Theorem 6 Let \mathcal{S}_k and H_k be (resp.) the complete structure and ‘big’ fact table of hypothesis k , and let Γ'_k be the repaired factorization of \mathcal{S}_k over H_k , and Y_0 the ‘explanation’ table where hypothesis k is recorded. Now, let \mathbf{Y}_k be a U-relational schema defined $\mathbf{Y}_k \triangleq \text{synthesize}(\mathcal{S}_k, H_k, Y_0)$. Then \mathbf{Y}_k is in BCNF w.r.t. Γ'_k and is minimal-cardinality.

Proof 15 We exploit the fact that the projection of $(\Gamma'_k)^+$ onto u-factor projections and predictive projections define a disjoint partition of $(\Gamma'_k)^+$ into its ϕ -projection $\Phi(\Gamma'_k)^+$ and v -projection $\Upsilon(\Gamma'_k)^+$. Since we know the form of fd's in each of them, the search space for BCNF violations is significantly reduced. The minimality of $|\mathbf{Y}_k|$ in turn comes from (Alg. 7) `merge`. See Appendix, §A.3.1. \square

5.5.2 Correctness of Uncertainty Decomposition

Recall from the preliminaries (cf. §5.1) that the U-relational equivalent of the relational product operation (main sub-operation of the join operation) has been introduced. Now, we provide the classical definition of a lossless join [20], i.e.,

when a decomposition of data from a relation into two or more relations is known to preserve the data in its original form by an application of the join.

Def. 22 Let $R[U]$ be a (U-)relational schema synthesized into collection $\mathbf{R} = \bigcup_{i=1}^m R_i$ and let Σ be an fd set on attributes U . We say that \mathbf{R} has a **lossless join** w.r.t. Σ if for every instance r of $R[U]$ satisfying Σ , we have $r = \bowtie_{i=1}^n \pi_{R_i}(r)$.

The lossless join property is of interest to ensure that our decomposition of the data from the ‘big’ fact table into u-factor projections preserves the data so that their join to ‘annotate’ the predictive projections when propagated by means of the U-relational join operation is correct. Theorem 7 guarantees that is the case.

Theorem 7 Let \mathcal{S}_k be the complete structure of hypothesis k , and $H_k[U]$ its ‘big’ fact table such that Γ'_k is the repaired factorization of \mathcal{S}_k over H_k and Y_0 is the ‘explanation’ table where hypothesis k is recorded. Now, let \mathbf{Y}_k be a U-relational schema defined $\mathbf{Y}_k \triangleq \text{synthesize}(\mathcal{S}_k, H_k, Y_0)$. Then,

- (a) the join $\bowtie_{i=1}^m Y_k^i [V_i D_i \mid \phi A_i G_i]$ of any subset of the u-factor projections of H_k is lossless w.r.t. Γ'_k .
- (b) any predictive projection $Y_k^j [\overline{V_j D_j} \mid S T]$, result of a join of the theoretical u-factor $Y_0 [V_0 D_0 \mid \phi v]$ with the ‘big’ fact table $H_k[U]$ and in turn with u-factor projections $Y_k^i [\overline{V_i D_i} \mid \phi A_i G_i]$, is lossless w.r.t. Γ'_k .

Proof 16 We make use of a lemma from Ullman [20, p. 397], and then the proof comes straightforwardly. See Appendix, §A.3.2. \square

Remark 8 The significance of Theorem 6 lies in that it guarantees the decomposition of uncertainty based on the causal ordering processing is in fact claim-centered as desirable for predictive analytics. Theorem 7 in turn is significant as it ensures that all the *empirical* uncertainty implicit in a hypothesis ‘big’ fact table H_k can be decomposed into u-factor projections that are (a) independent (not strongly correlated, cf. Problem 1), and (b) can be fully recovered by a join that is lossless w.r.t. repaired factorization Γ'_k of structure \mathcal{S}_k over H_k . This is essential to make sure that in u-propagation the composition of the required u-factors recovers the uncertainty associated with the predictive data. Since repaired factorization Γ'_k is

known to be a correct processing of the causal ordering (cf. results of Chapter 4), altogether Theorem 7 guarantees that the first causes are joined together correctly towards the predictive variables influenced by them. \square

As we have seen, the p-DB synthesis technique presented here is essentially targeted at design-theoretic properties. It is also motivated by computational performance, as uncertainty decomposition is desirable also to speed up probabilistic inference [17, p. 30-1]. In fact, the **U-intro** procedure is fully grounded in U-relations and p-WSA as implemented in the **MayBMS** system. Its computational performance is dominated by U-relational query processing. We present experimental studies on the **U-intro** procedure in §6.4, as they are designed from an applicability point of view. The goal is to provide some reference computational measures for prospective users.

5.6. Related Work

Informed on research on Graphical Models (GM) [69], Suciu et al. provide a striking motivation for work on *probabilistic database design* [17, p.30-1]. In GM design, probability distributions on large sets of random variables are decomposed into factors of simpler probability functions, over small sets of these variables. The factors can be identified, e.g., by using a set of axioms (the so-called ‘graphoids’) for reasoning about the probabilistic independence of variables [70]. The same design principle (sic.) applies to p-DB’s [17]: the data should be decomposed into its simplest components so that only key constraints hold in a table (i.e., it is in BCNF). Attribute- and tuple-level correlations should guide the table decomposition into simpler tables. Ideally, the original table with its probability distribution can be recovered as a query (a view) from the decomposed tables. We have followed such principle in our claim-centered decomposition for predictive analytics.

In fact, a connection between database normalization theory and factor decomposition in Graphical Models (GM) has been discussed by Verma and Pearl [70], but has not been explored since then. To date, there is no formal design theory for p-DB’s [17]. A step in that direction is taken by Sarma et al. [71]. Their initiative revisits dependency theory in view of reformulating fd’s for uncertain schema

design [71]. Our work takes a different direction. We refer to classical dependency theory and U-relational operations (viz, its uncertainty-introduction operator) to construct p-DB's systematically from scratch. We have focused on the extraction and processing of fd's towards a factorized U-relational schema. The synthesized schema is ensured to be in BCNF and have a lossless join.

Despite some major differences, our synthesis method builds upon the classical theory of relational schema design by synthesis [62]. Classical design by synthesis [62] was once criticized due to its too strong 'uniqueness' of fd's assumption [72, p. 443], as it reduces the problem of design to symbolic reasoning on fd's, arguably neglecting semantic issues. Probabilistic design, however, has roots in statistical design so that the problem is less amenable to human factors. As we extract the dependencies from a formal specification, design by synthesis is doing nothing but translating seamlessly (into fd's) the reduction made by the user herself in her tentative model for the studied phenomenon.

The last decade has seen significant research effort to make DB systems really usable [73]. Our design-by-synthesis framework can also be understood as a technique for *user-friendly p-DB design*. For instance, in comparison, the CRIUS system supports another kind of user-friendly DB design approach that provides users with a spreadsheet-like direct manipulation interface to increasingly add structure to their data [74]. Our dependency extraction and processing, instead, completely alleviates the user from the burden of data organization.

Also related to probabilistic DB design is the topic of conditioning a p-DB. It has been firstly addressed by Koch and Olteanu motivated by data cleaning applications [75]. They have introduced the **assert** operation to implement, as in AI, a kind of knowledge compilation, viz., world elimination in face of constraints (e.g., FDs). For hypothesis management, nonetheless, we need to apply *Bayes' conditioning* by asserting observed data, not constraints. In §2.5 we have presented an example that settles the kind of conditioning problem that is relevant to the Υ -DB vision. In Chapter 6 we present realistic use cases. We have addressed the problem at application level only in order to complete the realization of the vision in a real prototype system. The formulation of Bayes' conditioning as an extension of, say, the U-relational data model is open to future work (cf. §7.3).

5.7. Summary of Results

In this chapter we have studied and developed our end-purpose technique for the synthesis of a probabilistic DB geared for predictive analytics. It completes the pipeline (Fig. 1.4) so that conditioning can then be performed iteratively.

- Algorithm 8 `synthesize` gives a general formulation of how to perform uncertainty introduction from causal dependencies given in the form of fd's.
- By Problem 1, we have given a definition of uncertainty factor learning from data available in a given relation.
- Remark 7 provides an example of the u-factor and predictive projections resulting from p-DB synthesis and their corresponding probability distributions stored in the world table;
- By Remark 6 and Theorem 6, we have shown that U-relational schema \mathbf{Y}_k synthesized over the fd's processed by causal reasoning is in BCNF. That is, it is in fact a claim-centered decomposition as desirable for predictive analytics.
- Theorem 7 ensures that such (uncertainty) decomposition is correct, as the original (probability distribution) 'big' fact table is fully recoverable by a lossless join.

Chapter 6

Applicability

In this chapter we show the applicability of Υ -DB in real-world scenarios. We present use cases in Computational Physiology extracted from the Physiome project.¹ In §6.1 we introduce the Physiome project as providing a testbed for Υ -DB. Then in §6.2 we go through some Physiome case studies to show the construction of Υ -DB and its application for data-driven hypothesis management and analytics. In §6.3 we present a prototype of the Υ -DB system and demonstrate it through the running example introduced in §5.2. In §6.4 we present experiments on Physiome hypotheses. In §6.5 we provide a general discussion on the applicability of Υ -DB, its assumptions and scope. In §6.6 we conclude the chapter.

6.1. The Physiome Project as a Testbed

The Physiome project is an initiative to seriously address the problems of reproducibility, model integration and sharing in Computational Physiology [34, 33]. It essentially comprises:

- a curated *repository* of 380+ computational physiology models available online for researchers;²
- the *Mathematical Modeling Language* (MML) to allow models to be written in declarative form and then exported into a number of XML-compliant

¹ <http://physiome.org>.

² The Physiome model repository is expanded to over 73K+ models by including models extracted from other sources (such as the EBML-EBI BioModels DB, the CellML Archive, and the Kegg Pathways DB) and converted to MML automatically.

interoperable formats;³

- a *problem-solving environment* called **JSim** to allow researchers to code their MML models straightforwardly, run them under different parameter and solver settings and build customized data plots to see the results.

From the point of view of Υ -DB, Physiome is an external data source that provides a very interesting testbed with realistic scenarios. We extract Physiome models into Υ -DB by means of a wrapper we have implemented to read XMML files (JSim’s XML encoding of MML models). Simulation trial datasets are rendered by a parametrized UNIX script we have developed to invoke **JSim** automatically (in batch mode). Currently, Υ -DB’s Physiome wrapper is designed to read MAT files to load both the model input (parameter settings data) and its associated model output (computed predictive data) for each simulation trial.

Physiome does not keep records of phenomena in a repository, but it does have observational data attached to some of the entries of the model repository. Such models appear in the filter ‘models with data,’ meaning that they have one or more observational datasets and plots showing how the model data fits to observations. We shall make use of model entries containing observational data in the realistic scenarios presented in this paper.

6.2. Case Studies

In this section we present use cases extracted from the Physiome model repository.⁴

6.2.1 Case: Hemoglobin Oxygen Saturation

In this case we stress the potential of data-driven hypothesis analytics in comparison to handcrafted curve fitting (visual) analysis. We study three different hypotheses that perform “closely” visually when compared to their target phenomenon dataset, see Fig. 6.1. All of them have been empirically set as fit as possible to the observations (‘R1s1’ dataset) in their local view (in separate), and are now compared together in a global view.

³ http://www.physiome.org/jsim/docs/MML_Intro.html.

⁴ <http://www.physiome.org/Models/modelDB/>.

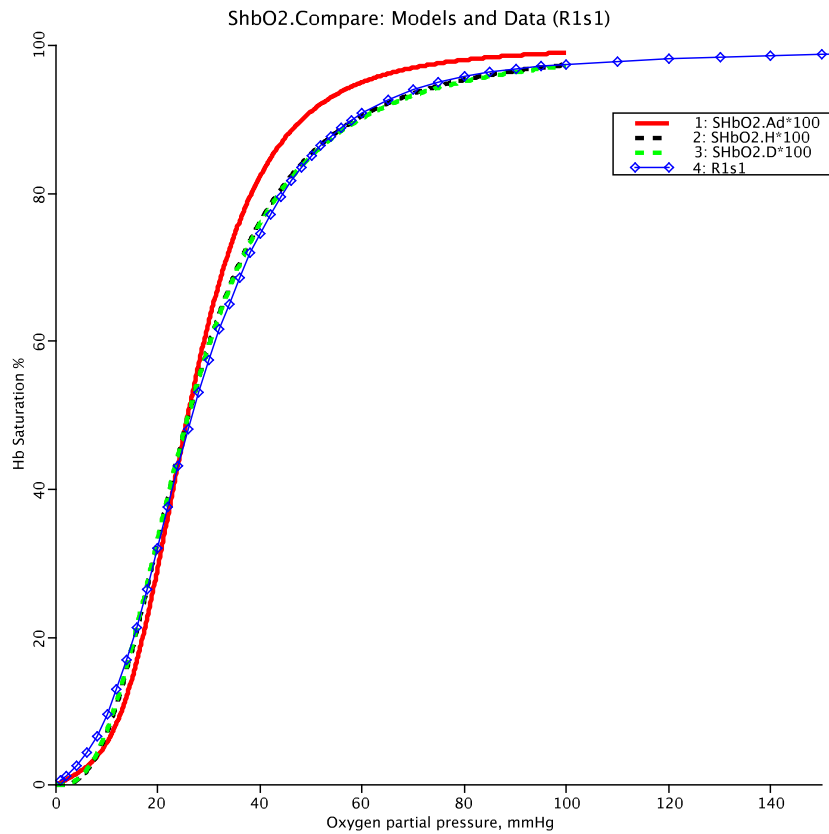


Figure 6.1. Plot of hemoglobin oxygen saturation hypotheses (SHbO2.{Ad, H, D} curves) and their target observations ('R1s1' dataset). (source: Physiome).

Example 8 The resources of this example are shown in Fig. 6.2. We consider the Physiome model entries described in relation HYPOTHESIS, associated to the phenomenon described in relation PHENOMENON (cf. explanation relation H_0). One single hypothesis trial (its best fit) is considered for each hypothesis. \square

HYPOTHESIS	v	Name	Description
	28	HbO.Hill	Hill Equation for O2 binding to hemoglobin.
	31	HbO.Adair	Hemoglobin O2 saturation curve using Adair's 4-site equation.
	32	HbO.Dash	Hemoglobin O2 saturation curve at varied levels of PCO2 and pH.

PHENOMENON	ϕ	Description
	1	Hemoglobin oxygen saturation with observational dataset from Sevenringhaus 1979.

H_0	ϕ	v
	1	28
	1	31
	1	32

Figure 6.2. Descriptive (textual) data of Example 8, with ids v from Physiome's model repository (<http://www.physiome.org/Models/modelDB/>).

Encoding. The fd encoding of hypotheses $v \in \{28, 31, 32\}$ is shown (resp.) in Fig. 6.3, Fig. 6.4 and Fig. 6.5.

$$\Sigma_{28} = \{ \begin{array}{l} \text{KO2 } n \text{ pO2 } v \rightarrow \text{SHbO2_H}, \\ \phantom{\text{KO2 }} n \text{ p50 } v \rightarrow \text{KO2}, \\ \phi \rightarrow n \text{ p50 pO2_delta pO2_max pO2_min } \end{array} \}.$$

Figure 6.3. Fd set Σ_{28} of hypothesis $v=28$.

$$\Sigma_{31} = \{ \begin{array}{l} \text{A1 A2 A3 A4 pO2 } v \rightarrow \text{SHbO2_Ad}, \\ \phantom{\text{A1 A2 A3 A4 }} p50 v \rightarrow \text{A1 A2 A4}, \\ \phi \rightarrow \text{A3 p50 pO2_delta pO2_max pO2_min } \end{array} \}.$$

Figure 6.4. Fd set Σ_{31} of hypothesis $v=31$.

$$\Sigma_{32} = \{ \begin{array}{l} \text{KO2 pO2 } v \rightarrow \text{SHbO2_D}, \\ \phantom{\text{KO2 }} \text{term1 term2 term3 } v \rightarrow \text{KO2}, \\ \phantom{\text{KO2 }} \text{HpRBC K6p alphaO2 } v \rightarrow \text{term1}, \\ \text{HpRBC K3 K5 alphaCO2 pCO2 } v \rightarrow \text{term2}, \\ \text{HpRBC K2 K4 alphaCO2 pCO2 } v \rightarrow \text{term3}, \\ \phantom{\text{KO2 }} \text{pH } v \rightarrow \text{HpRBC}, \\ \phi \rightarrow \text{K2 K3 K4 K5 K6p alphaCO2 alphaO2 pCO2 pH pO2_delta, pO2_max pO2_min } \end{array} \}.$$

Figure 6.5. Fd set Σ_{32} of hypothesis $v=32$.

Symbol Mappings. As we have seen, the insertion of hypothesis trial datasets requires users to specify a target phenomenon and the corresponding mappings from the hypothesis symbols to the target phenomenon symbols. In this use case, we have:

- $\mathcal{M}_{28 \mapsto 1} = \{ \text{pO2} \mapsto \text{pO2}, \text{SHbO2_H} \mapsto \text{SHbO2} \};$
- $\mathcal{M}_{31 \mapsto 1} = \{ \text{pO2} \mapsto \text{pO2}, \text{SHbO2_Ad} \mapsto \text{SHbO2} \};$
- $\mathcal{M}_{32 \mapsto 1} = \{ \text{pO2} \mapsto \text{pO2}, \text{SHbO2_D} \mapsto \text{SHbO2} \};$

Hypothesis Management. Query Q1 illustrates the feature of hypothesis management for this case. We consider the user is interested in all SHbO2 predictions over a subset of the pO2 domain. Its result set is shown in Fig. 6.6.

Q1. (select phi, epsilon, tid, "pO2", "SHbO2_H" as SHbO2 from Y28_claim1
 where phi=1 and "pO2">=20 and "pO2"<=40) union all
 (select phi, epsilon, tid, "pO2", "SHbO2_Ad" as SHbO2 from Y31_claim1
 where phi=1 and "pO2">=20 and "pO2"<=40) union all
 (select phi, epsilon, tid, "pO2", "SHbO2_D" as SHbO2 from Y32_claim1
 where phi=1 and "pO2">=20 and "pO2"<=40) order by "pO2", epsilon, tid;

Q1	ϕ	v	tid	pO2	SHbO2
	1	28	1	20	0.329956122828398
	1	31	1	20	0.294443723056007
	1	32	1	20	0.334165672301096

	1	28	1	40	0.761898061367189
	1	31	1	40	0.823463829100424
	1	32	1	40	0.759042287799556

Figure 6.6. Result set of hypothesis management query Q1.

Hypothesis Analytics. Fig. 6.7 shows the results of analytics after conditioning the probability distribution in the presence of observations ('R1s1' dataset). The fact that hypothesis $v = 31$ provides the best explanation for the studied phenomenon is enabled by the application of Bayesian inference as implemented within the Υ -DB system. The contribution of the Υ -DB methodology is to equip users with a tool for large-scale, data-driven hypothesis management and analytics.

PH1_CONF	ϕ	v	tid	pO2	SHbO2	Prior	Posterior
	1	28	1	1	0.000151184162020125	.333	.326
	1	31	1	1	0.003789100566457180	.333	.349
	1	32	1	1	0.000178973375779681	.333	.325

	1	28	1	100	0.974346796798538	.333	.326
	1	31	1	100	0.990781330988763	.333	.349
	1	32	1	100	0.972764121981342	.333	.325

Figure 6.7. Results of analytical study on the hemoglobin phenomenon.

6.2.2 Case: Baroreflex Dysfunction in Dahl SS Rat

This case is extracted from the Virtual Physiological Rat project,⁵ Here we show the potential of data-driven hypothesis management and analytics for model tuning. Fig. 6.8 shows the best fit of a baroreflex model for an observational dataset acquired by experiment on Dahl SS rat [76]. We in turn use Υ -DB to carry out such hypothesis management and analytics. We generate by a parameter sweep script 1K trials and insert them into the database. A best fit is then selected automatically by Bayesian inference.

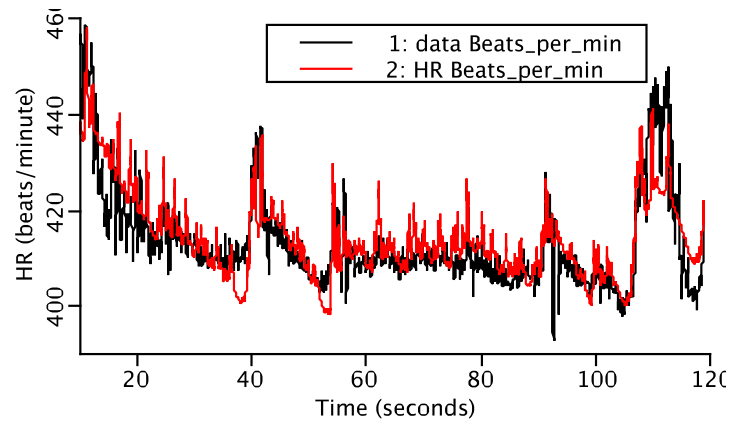


Figure 6.8. Plot of baroreflex hypothesis ('HR') for Dahl SS Rat and its target observations ('data'). (source: [76]).

Example 9 The resources of this example are shown in Fig. 6.9. We consider the single hypothesis entry described in relation HYPOTHESIS, and the phenomenon described in relation PHENOMENON. By parameter sweep, 1K trials are inserted into Υ -DB for management and analytics. \square

HYPOTHESIS	v	Name	Description
	1001	Baroreflex_SB_CT	Physiological model of the full baroreflex heart control system based on experimental measurements.

PHENOMENON	ϕ	Description
	2	Baroreflex dysfunction in Dahl SS Rat.

H_0	ϕ	v
	2	1001

Figure 6.9. Descriptive (textual) data of Example 9.

Encoding. The fd encoding of hypothesis $v \in \{1001\}$ is shown in Fig. 6.12.

⁵ <http://virtualrat.org/computational-models/>.

Symbol Mappings. We consider that the user provides symbol mappings:

- $\mathcal{M}_{1001 \mapsto 2} = \{ \text{Time} \mapsto \text{Time}, \text{HR} \mapsto \text{HR} \};$

Hypothesis Management. In query Q2 we consider that the user is interested in time instants where the heart rate is higher than a threshold, say, 300 beats/min. The result set is shown in Fig. 6.10.

Q2. **select** phi, upsilon, tid, “Time”, “HR” **from** Y1001_claim1
where phi=2 **and** “HR”>=300 **order by** “Time”, tid;

Q2	ϕ	v	tid	Time	HR
	2	1001	1	0.61	300.013659905941
	2	1001	2	0.61	300.011268345391

	2	1001	96	0.61	300.001934440349
	2	1001	1	0.62	300.607671377207

Figure 6.10. Result set of hypothesis management query Q2.

Hypothesis Analytics. Fig. 6.11 shows the results of analytics on phenomenon $\phi=2$ after conditioning the probability distribution in the presence of observations (‘SSBN9_HR’ dataset). Since this case deals with model tuning, viz., 1K slightly different parameter settings, the trial ranking is decided by small differences in the posterior probability distribution (cf. Fig 6.11).

PH2_CONF	ϕ	v	tid	Time	HR	Prior	Posterior
	2	1001	491	0.4	286.556506432110	.001000	.00103159
	2	1001	591	0.4	286.525209765226	.001000	.00103144
	2	1001	492	0.4	286.555565558231	.001000	.00103023

	2	1001	491	118.9	421.251853783050	.001000	.00103159
	2	1001	591	118.9	421.110425308905	.001000	.00103144
	2	1001	492	118.9	421.297317710657	.001000	.00103023

Figure 6.11. Results of analytical study on the baroreflex phenomenon.

6.2.3 Case: Myogenic Behavior of a Blood Vessel

Computational models of physiology may account for diverse effects that take place at different levels of biological organization from the organ to the cellular

$$\Sigma_{1001} = \{$$

$$\begin{aligned}
& \text{HR } v \rightarrow \text{Period}, \\
& \text{Beta HR}_p \text{ HR}_s \text{ HR}_{\min} \text{ HR}_o v \rightarrow \text{HR}, \\
& \text{HR}_o \text{ delta_HR}_p v \rightarrow \text{HR}_p, \\
& \text{delta_HR}_{\text{pfast}} \text{ delta_HR}_{\text{pslow}} v \rightarrow \text{delta_HR}_p, \\
& \text{HR}_o \text{ delta_HR}_s v \rightarrow \text{HR}_s, \\
& \text{Time delta_HR}_s \text{ Time}_{\min} \text{ delta_HR}_{ss} \text{ tau_HR}_{\text{nor}} v \rightarrow \text{delta_HR}_s, \\
& \text{Gamma delta_HR}_{ps} v \rightarrow \text{delta_HR}_{\text{pfast}}, \\
& \text{Gamma Time delta_HR}_{ps} \text{ delta_HR}_{\text{pslow}} \text{ Time}_{\min} \text{ tau_HR}_{\text{ach}} v \rightarrow \text{delta_HR}_{\text{pslow}}, \\
& \text{K}_{\text{nor}} \text{ c}_{\text{nor}} \text{ delta_HR}_{s\text{max}} v \rightarrow \text{delta_HR}_{ss}, \\
& \text{C}_{\text{ach}} \text{ K}_{\text{ach}} \text{ delta_HR}_{p\text{max}} v \rightarrow \text{delta_HR}_{ps}, \\
& \text{Time Ts c}_{\text{nor}} \text{ Time}_{\min} \text{ q}_{\text{nor}} \text{ tau}_{\text{nor}} v \rightarrow \text{c}_{\text{nor}}, \\
& \text{C}_{\text{ach}} \text{ Time}_{\min} \text{ Time Tp q}_{\text{ach}} \text{ tau}_{\text{ach}} v \rightarrow \text{C}_{\text{ach}}, \\
& \text{Gs T}_{\text{smax}} \text{ T}_{\text{smin}} \text{ alpha}_{\text{cns}} \text{ alpha}_{s0} v \rightarrow \text{Ts}, \\
& \text{Gp T}_{p\text{max}} \text{ T}_{p\text{min}} \text{ alpha}_{\text{cns}} \text{ alpha}_{p0} v \rightarrow \text{Tp}, \\
& \text{Gcns n } v \rightarrow \text{alpha}_{\text{cns}}, \\
& \text{S Zeta delta delta}_{\text{th}} v \rightarrow \text{n}, \\
& \text{Eps}_1 \text{ Eps}_{\text{wall}} v \rightarrow \text{delta}, \\
& \text{B1 Eps}_1 \text{ Time}_{\min} \text{ Eps}_2 \text{ Eps}_2 \text{ Time}_{\min} \text{ Eps}_{\text{wall}} \text{ K1 Kne Time } v \rightarrow \text{Eps}_1, \\
& \text{B1 B2 Eps}_1 \text{ Eps}_1 \text{ Time}_{\min} \text{ Eps}_2 \text{ Time}_{\min} \text{ Eps}_3 \text{ Eps}_3 \text{ Time}_{\min} \text{ K1 K2 Time } v \rightarrow \text{Eps}_2, \\
& \text{B2 B3 Eps}_2 \text{ Eps}_2 \text{ Time}_{\min} \text{ Eps}_3 \text{ Time}_{\min} \text{ K2 K3 Time } v \rightarrow \text{Eps}_3, \\
& \text{R R0 } v \rightarrow \text{Eps}_{\text{wall}}, \\
& \text{A } v \rightarrow \text{R}, \\
& \text{A Time}_{\min} \text{ B}_{\text{wall}} \text{ C}_{\text{wall}} \text{ P R0 Time } v \rightarrow \text{A}, \\
& \text{HR}_{\text{max}} \text{ HR}_o v \rightarrow \text{delta_HR}_{s\text{max}}, \\
& \text{HR}_{\min} \text{ HR}_o v \rightarrow \text{delta_HR}_{p\text{max}}, \\
& \text{Time } \phi \rightarrow \text{data}, \\
& \phi \rightarrow \text{A Time}_{\min} \text{ B1 B2 B3 Beta B}_{\text{wall}} \text{ C}_{\text{ach}} \text{ Time}_{\min} \text{ C}_{\text{wall}} \text{ Eps}_1 \text{ Time}_{\min} \\
& \text{Eps}_2 \text{ Time}_{\min} \text{ Eps}_3 \text{ Time}_{\min} \text{ Gamma Gcns Gp Gs HR}_{\text{max}} \text{ HR}_{\min} \text{ HR}_o \text{ K1 K2 K3} \\
& \text{K}_{\text{ach}} \text{ K}_{\text{nor}} \text{ Kne P R0 S Time}_{\text{delta}} \text{ Time}_{\text{max}} \text{ Time}_{\min} \text{ T}_{p\text{max}} \text{ T}_{p\text{min}} \text{ T}_{s\text{max}} \text{ T}_{s\text{min}} \\
& \text{Zeta alpha}_{p0} \text{ alpha}_{s0} \text{ c}_{\text{nor}} \text{ Time}_{\min} \text{ delta_HR}_{\text{pslow}} \text{ Time}_{\min} \text{ delta_HR}_s \text{ Time}_{\min} \\
& \text{delta}_{\text{th}} \text{ q}_{\text{ach}} \text{ q}_{\text{nor}} \text{ tau_HR}_{\text{ach}} \text{ tau_HR}_{\text{nor}} \text{ tau}_{\text{ach}} \text{ tau}_{\text{nor}} \}.
\end{aligned}$$

Figure 6.12. Fd set Σ_{1001} of hypothesis $v=1001$.

and molecular levels [33]. Typically, a sophisticated model is developed incrementally by, say, adding detail into some previously existing model or extending its dimensionality (e.g., extending it from a stationary to a dynamic account of phenomena). In this case study (cf. Example 10) we consider alternative models of the myogenic behavior of a reference human blood vessel.

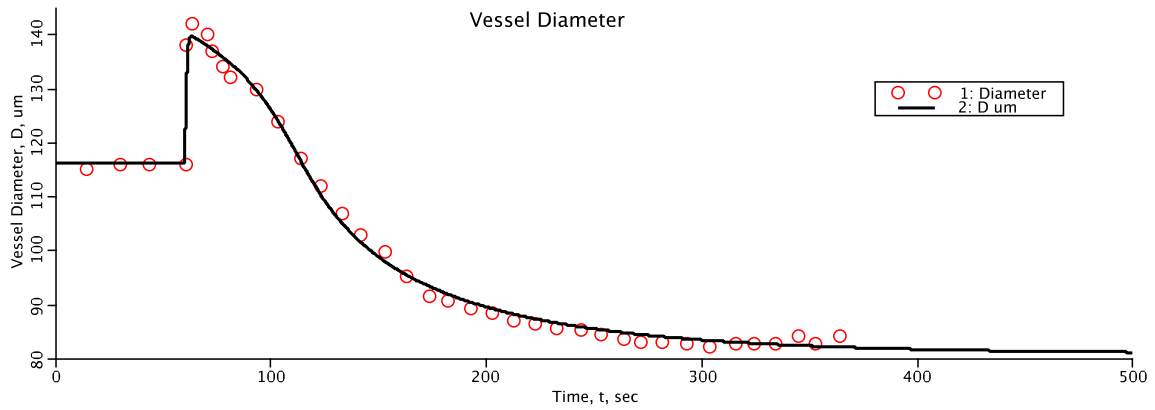


Figure 6.13. Plot of myogenic behavior hypothesis ('D') according to $v=89$, trial $\text{tid}=2$, and its target observations ('Diameter').

Example 10 (See Fig. 6.14). We consider the Physiome model entries displayed in relation HYPOTHESIS, and two phenomena (see relation PHENOMENON). One trial is considered for hypothesis $v = 60$, and two for hypothesis $v = 89$. \square

HYPOTHESIS	v	Name	Description
	60	Myogenic_Compliant_Vessel	This model simulates the flow through a passive and actively responding vessel driven by a sinusoidal pressure input.
	89	Myo_Dyn_Resp_wFit	This model describes the dynamic response of a vessel after a step increase in intraluminal pressure.

PHENOMENON	ϕ	Description
	3	Dynamics of vessel diameter in response to pulsatile intraluminal pressure.

H_0	ϕ	v
	3	60
	3	89

Figure 6.14. Descriptive (textual) data of Example 10.

Encoding. The fd encoding of hypotheses $v \in \{60, 89\}$ is shown (resp.) in Fig. 6.15 and Fig. 6.16.

Symbol Mappings. We consider that the user provides symbol mappings:

- $\mathcal{M}_{60 \rightarrow 1} = \{ t \mapsto \text{Time}, D \mapsto \text{Diameter} \};$
- $\mathcal{M}_{89 \rightarrow 1} = \{ t \mapsto \text{Time}, D \mapsto \text{Diameter} \};$

$$\begin{aligned}
\Sigma_{60} = \{ & \text{Fcomp Fout } v \rightarrow \text{Fin}, \\
& \text{Pin Pout R } v \rightarrow \text{Fout}, \\
& \text{V V_t_min t } v \rightarrow \text{Fcomp}, \\
& \text{D L mu } v \rightarrow \text{R}, \\
& \text{D L } v \rightarrow \text{V}, \\
& \text{D Pin T } v \rightarrow \text{A}, \\
& \text{A C1a C1p C2a C2p C3a Dp100 Pin T Ttarget } v \rightarrow \text{Atarget}, \\
& \text{Atarget Cglobal Cmyo Pin } v \rightarrow \text{D}, \\
& \text{D D_t_min Dc Tc Ttarget t taud } v \rightarrow \text{T}, \\
& \text{A A_t_min Atarget t taua } v \rightarrow \text{Ttarget}, \\
& \text{Ac } v \rightarrow \text{A_t_min}, \\
& \text{Cglobal Cmyo Tc } v \rightarrow \text{Ac}, \\
& \text{Dc } v \rightarrow \text{D_t_min}, \\
& \text{Dc Pmean } v \rightarrow \text{Tc}, \\
& \text{Pamp Pmean t tnorm } v \rightarrow \text{Pin}, \\
& \text{C1a C1p C2a C2p C3a Cglobal Cmyo Dp100 Pc } v \rightarrow \text{Dc}, \\
& \phi \rightarrow \text{C1a C1p C2a C2p C3a Cglobal Cmyo Dp100 L Pamp Pc Pext} \\
& \text{Pmean Pout V_t_min mu t_delta t_max t_min taua taud tnorm } \}.
\end{aligned}$$

Figure 6.15. Fd set Σ_{60} of hypothesis $v=60$.

$$\begin{aligned}
\Sigma_{89} = \{ & \text{D P T } v \rightarrow \text{A}, \\
& \text{A C1a C1p C2a C2p C3a Dp100 P T Ttarget } v \rightarrow \text{Atarget}, \\
& \text{Atarget Cglobal Cmyo P } v \rightarrow \text{D}, \\
& \text{D D_t_min Dc Tc Ttarget t taud } v \rightarrow \text{T}, \\
& \text{A A_t_min Atarget t taua } v \rightarrow \text{Ttarget}, \\
& \text{Ac } v \rightarrow \text{A_t_min}, \\
& \text{Dc Pc } v \rightarrow \text{Tc}, \\
& \text{Cglobal Cmyo Dc Pc } v \rightarrow \text{Ac}, \\
& \text{Dc } v \rightarrow \text{D_t_min}, \\
& \text{DelP Pc } v \rightarrow \text{P}, \\
& \text{C1a C1p C2a C2p C3a Cglobal Cmyo Dp100 Pc } v \rightarrow \text{Dc}, \\
& \text{t } \phi \rightarrow \text{DelP}, \\
& \phi \rightarrow \text{C1a C1p C2a C2p C3a Cglobal Cmyo Dp100 Pc t_delta t_max t_min taua taud } \}.
\end{aligned}$$

Figure 6.16. Fd set Σ_{89} of hypothesis $v=89$.

Hypothesis Management. Query Q3 illustrates the feature of hypothesis management for this case. The user selects all diameter predictions within the time interval $t \in [100, 300]$ (cf. plot in Fig. 6.13). Its result set is shown in Fig. 6.17.

Q3. **select** phi, epsilon, tid, "t", "D" **from** Y60_claim1
where phi=3 **and** "t">=100 **and** "t"<=300 **union all**
select phi, epsilon, tid, "t", "D" **from** Y89_claim1
where phi=3 **and** "t">=100 **and** "t"<=300 **order by** "t", epsilon, tid;

Q3	ϕ	v	tid	t	D
	3	89	1	100.00	194.622865847211
	3	89	1	100.00	97.3787340059609
	3	89	2	100.00	126.167727083098
	3	89	1	100.01	194.626017703936
	3	89	1	100.01	98.0174705905828
	3	89	2	100.01	126.161751822302

Figure 6.17. Result set of hypothesis management query Q3.

Hypothesis Analytics. Fig. 6.18 shows the results of analytics on phenomenon $\phi = 3$ after conditioning the probability distribution in the presence of observations, viz., 'Davis_Sikes_Fig3_Myo_DigData' dataset.

PH3_CONF	ϕ	v	tid	Time	Diameter	Prior	Posterior
	3	60	1	14.8	194.996792066637	.5	.000
	3	89	1	14.8	97.0568250956827	.25	.269
	3	89	2	14.8	116.327813203282	.25	.731

	3	60	1	30.5	195.684170988267	.5	.000
	3	89	1	30.5	97.0568250767575	.25	.269
	3	89	2	30.5	116.327813337087	.25	.731

Figure 6.18. Results of analytics on the vessel's myogenic behavior phenomenon.

In this case study two tentative models have been considered under a uniform prior probability distribution which has been updated to a posterior distribution. Note that, even though hypothesis $v = 60$ has its probability weight concentrated in a single trial, the Bayesian inference is able to indicate $v = 89$ as the best explanation for $\phi = 3$ and tid = 2, in particular, its best fit.

6.3. System Prototype

A first prototype of the Υ -DB system has been implemented as a Java web application, with the pipeline component in the server side on top of MayBMS (a backend

extension of PostgreSQL). We have developed a demonstration of this prototype (cf. [28]), in which we go through the whole design-by-synthesis pipeline (Fig. 1.4) exploring use case scenarios. In this section we provide a brief demonstration of the system in the population dynamics scenario previously introduced in this thesis.

The demonstration unfolds in three phases. In the first phase, we show the ETL process to give a sense of what the user has to do in terms of simple phenomena description, hypothesis naming and file upload to get her phenomena and hypotheses available in the system to be managed as data. In the second phase, we reproduce some typical queries of hypothesis management (like those shown in the previous section). In the third phase, we enter the hypothesis analytics module. The user chooses a phenomenon for a hypothesis evaluation study, and the system lists all the predictions with their probabilities under some selectivity criteria (e.g., population at year 1920). The predictions are ranked according to their probabilities, which are conditioned on the observational data available for the chosen phenomenon.

6.3.1 Demo Screenshots

Fig. 6.21 shows screenshots of the system. Fig. 6.21(a) shows the research projects currently available for a user. Figs. 6.21(b, c) show the ETL interfaces for phenomenon and hypothesis data definition (by synthesis), and then the insertion of hypothesis trial datasets, i.e., explanations of a hypothesis towards a target phenomenon. Fig. 6.21(d) shows the interface for basic hypothesis management by listing the predictions of a given simulation trial. Figs. 6.21(e, f) show two tabs of the hypothesis analytics module, viz., selection of observations and then viewing the corresponding alternative predictions ranked by their conditioned probabilities.

6.3.2 Demo Case: Population Dynamics

In this case we refer to a well-known problem in Computational Science, viz., population dynamics scenarios, to demonstrate the Υ -DB system prototype. Fig. 6.19 shows census data collected from in the US from 1790 to 1990.⁶ Fig. 6.20 shows observational data collected from Hudson’s Bay from 1900 to 1920 on the Lynx-Hare population [77].

⁶ Cf. <https://www.census.gov/population/censusdata/table-4.pdf>.

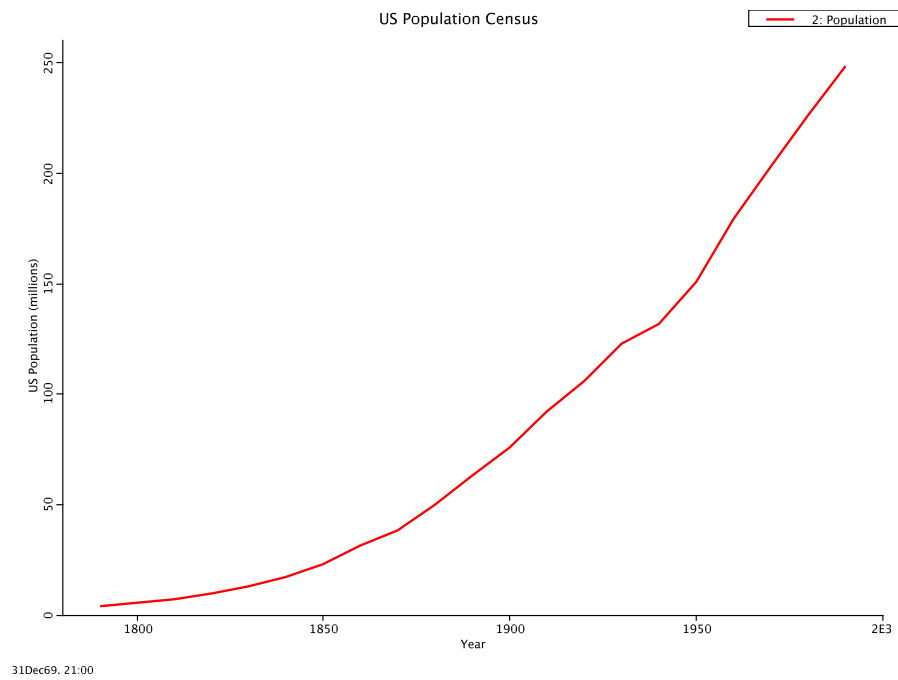


Figure 6.19. Census US population from 1790 to 1990.

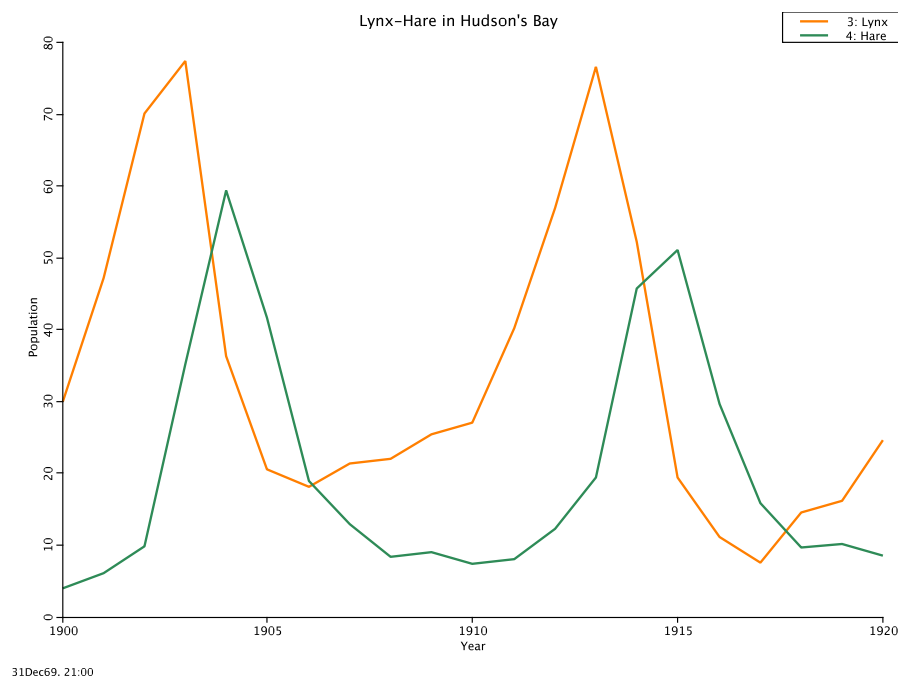
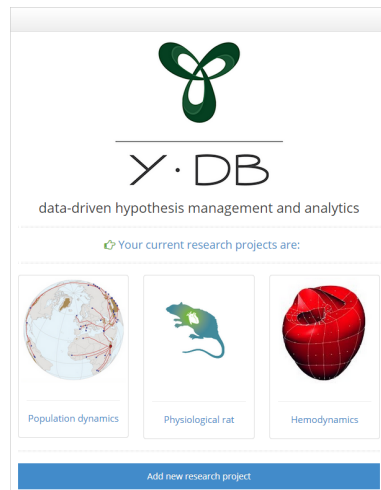


Figure 6.20. Lynx-Hare population observed in Hudon's Bay from 1900 to 1920.



(a) Research dashboard after login.

Phenomenon data definition

Phenomenon id
2

Research
Population dynamics

Description
Lynx population in Hudson's Bay, Canada, from 1900 to 1920.

Upload dataset (CSV format)
[Choose File](#) Lynx_Hare.csv

Loading observations
70%

Observable

Year

Lynx

(b) Phenomenon data definition.

Hypothesis data definition

Hypothesis id
3

Name
Lotka-Volterra model

Upload structure (XML format)
[Choose File](#) Lotka_Volterra.xml

Hypothesis structure: 100%

Phenomenon
Lynx population in Hudson's Bay, Canada, from 1900 to 1920.

Map symbols

Variable	Observable
t	Year
x	Lynx

Hypothesis trial datasets (MAT format)
[Choose Files](#) 10 files

(c) Hypothesis data definition.

Hypothesis management

Hypothesis
Lotka-Volterra model

Phenomenon
Lynx population in Hudson's Bay, Canada, from 1900 to 1920.

Simulation trial
6

Key 1 Key 2

t	y	x
1904.0	31.1083920070696	75.9196961932191
1904.1	34.1895828779035	74.4878105315043
1904.2	37.4356490187431	72.6675156977604
1904.3	40.8008312965646	70.4705874385644
1904.4	44.2262696820135	67.9230251780413
1904.5	47.6417755935588	65.0649920502872
1904.6	50.9691309633260	61.9493743700121
1904.7	54.1267031673551	58.6389813583626

(d) Hypothesis management.

Hypothesis Analytics

Phenomenon
Lynx population in Hudson's Bay, Canada, from 1900 to 1920.

Observations Predictions

	Year	Lynx	Hare
✓	1900	30	4
✓	1901	47.2	6.1
✓	1902	70.2	9.8
✓	1903	77.4	35.2
✓	1904	36.3	59.4
✓	1905	20.6	41.7
✓	1906	18.1	19
✓	1907	21.4	13
✓	1908	22	8.3

(e) Analytics: selected observations tab.

Hypothesis Analytics

Phenomenon
Lynx population in Hudson's Bay, Canada, from 1900 to 1920.

Observations Predictions

upsilon	tid	Year	Lynx	conf
3	2	1904	65.060410460081	0.183505
3	6	1904	75.919696193219	0.179993
3	4	1904	77.459735769215	0.175992
3	1	1904	89.592307430943	0.131452
3	5	1904	88.321831841064	0.127000
3	3	1904	90.083803232660	0.124023
1	1	1904	16.487212706992	0.047211
2	2	1904	77.822475573932	0.017372
2	1	1904	79.812581025093	0.013234
1	2	1904	18.221188003898	0.000220

(f) Analytics: ranked predictions tab.

Figure 6.21. Screenshots of this first prototype of the Υ -DB system.

Example 11 (See Fig. 6.22). We consider the model entries displayed in relation HYPOTHESIS, and two phenomena (see relation PHENOMENON). For $\phi = 1$, three trials are considered for hypothesis $v = 1$ and six for hypothesis $v = 2$. For $\phi = 2$, in turn, two trials are considered for hypothesis $v = 1$ and $v = 2$, and six trials for hypothesis $v = 3$. Note the data definition interfaces in Figs. 6.21(b, c). \square

HYPOTHESIS	v	Name	Description
	1	Malthusian growth model	Exponential growth model 'growth in population is proportional to its size' is considered the first principle of population dynamics.
	2	Logistic equation	This model introduces growth saturation to the Malthusian model due to the limitation of resources.
	3	Lotka-Volterra model	This model describes predator-prey interactions.

PHENOMENON	ϕ	Description
	1	US population from 1790 to 1990.
	2	Lynx population in Hudson's Bay, Canada, from 1900 to 1920.

H_0	ϕ	v
	1	1
	1	2
	2	1
	2	2
	2	3

Figure 6.22. Descriptive (textual) data of Example 11.

Encoding. The fd encoding of hypotheses $v \in \{1, 2, 3\}$ is shown (resp.) in Fig. 6.23, Fig. 6.24 and Fig. 6.25. See hypothesis structure processing in Fig. 6.21(c).

$$\Sigma_1 = \{ \text{b t x_t_min } v \rightarrow x, \\ \phi \rightarrow \text{b t_delta t_max t_min x_t_min} \}.$$

Figure 6.23. Fd set Σ_1 of hypothesis $v = 1$.

$$\Sigma_2 = \{ \text{K b t x_t_min } v \rightarrow x, \\ \phi \rightarrow \text{K b t_delta t_max t_min x_t_min} \}.$$

Figure 6.24. Fd set Σ_2 of hypothesis $v = 2$.

Symbol Mappings. We consider that the user provides the following symbol mappings for (resp.) phenomena $\phi = 1$ and $\phi = 2$, see the interface for mapping symbols in Fig. 6.21(c).

$$\Sigma_3 = \{ \begin{array}{l} \mathbf{b} \ \mathbf{p} \ \mathbf{t} \ \mathbf{y} \ \mathbf{v} \rightarrow \mathbf{x}, \\ \mathbf{d} \ \mathbf{r} \ \mathbf{t} \ \mathbf{x} \ \mathbf{v} \rightarrow \mathbf{y}, \\ \phi \rightarrow \mathbf{b} \ \mathbf{d} \ \mathbf{p} \ \mathbf{r} \ \mathbf{t_delta} \ \mathbf{t_max} \ \mathbf{t_min} \ \mathbf{x_t_min} \ \mathbf{y_t_min} \end{array} \}.$$

Figure 6.25. Fd set Σ_3 of hypothesis $v=3$.

- $\mathcal{M}_{1 \mapsto 1} = \{ t \mapsto \text{Year}, x \mapsto \text{Population} \};$
- $\mathcal{M}_{2 \mapsto 1} = \{ t \mapsto \text{Year}, x \mapsto \text{Population} \};$
- $\mathcal{M}_{1 \mapsto 2} = \{ t \mapsto \text{Year}, x \mapsto \text{Lynx} \};$
- $\mathcal{M}_{2 \mapsto 2} = \{ t \mapsto \text{Year}, x \mapsto \text{Lynx} \};$
- $\mathcal{M}_{3 \mapsto 2} = \{ t \mapsto \text{Year}, x \mapsto \text{Lynx} \};$

Hypothesis Management. Query Q4 illustrates the feature of hypothesis management for this case. The user selects hypothesis $v=3$ (the Lotka-Volterra model), and filters its available data for trial $\text{tid}=6$ on phenomenon $\phi=2$. Both the form-based query set-up and its result set are shown in Fig. 6.21(d).

Q4. select “t”, “y”, “x” from Y3_claim1
where $\text{upsilon}=3$ **and** $\text{phi}=2$ **and** $\text{tid}=6$ **order by** “t”;

Hypothesis Analytics. Fig. 6.26 and Fig. 6.27 show the results of analytics on (resp.) phenomena $\phi=1$ and $\phi=2$ after conditioning the probability distribution in the presence of (resp.) observational datasets ‘US-census’ and ‘Lynx-Hare.’ In the first one, the user verifies that hypothesis $v=1$ (the Malthusian model) is unlikely to be competitive with hypothesis $v=2$ (the Logistic equation) as an approximation of the US population dynamics from 1790 to 1990. That is, if the user knows her current trials are reasonable, then more trials on the Malthusian model hardly could outperform trials on the Logistic equation for the studied phenomenon.

PH1_CONF	ϕ	v	tid	Year	Population	Prior	Posterior
...
1	1	1	1	1920	194.102222534948	.250	.000000
1	1	2	2	1920	140.244165184248	.250	.000000
1	2	1	1	1920	82.3031951115155	.125	.133038
1	2	2	2	1920	108.251924734215	.125	.239684
1	2	3	3	1920	105.918217777077	.125	.290026
1	2	4	4	1920	105.988231944275	.125	.337251
...

Figure 6.26. Results of analytics on the US population phenomenon.

PH2_CONF	ϕ	v	tid	Year	Lynx	Prior	Posterior
...
2	1	1	1	1904	16.49	.167	.047
2	1	2	2	1904	18.22	.167	.000
2	2	1	1	1904	79.81	.167	.013
2	2	2	2	1904	77.82	.167	.017
2	3	1	1	1904	89.59	.055	.131
2	3	2	2	1904	65.06	.055	.184
2	3	3	3	1904	90.08	.055	.124
2	3	4	4	1904	77.46	.055	.176
2	3	5	5	1904	88.32	.055	.127
2	3	6	6	1904	75.92	.055	.180
...

Figure 6.27. Results of analytics on the Hudson’s Bay lynx population phenomenon, see interfaces in Figs. 6.21(e, f).

6.4. Experiments

The efficiency and scalability of the U-relational representation system and its p-WSA query algebra have been extensively demonstrated [27]. Υ -DB’s, as U-relational hypothesis DB’s, must therefore be as efficient and scalable as any arbitrary U-relational DB.

In these experiments (see Fig. 6.28) we provide some measures of performance of the method of Υ -DB in the particular context of our real-world Physioime testbed. Our purpose here is to provide a concrete feel on how efficient the Υ -DB methodology can be. However, most of these tests (the four graphs on the bottom in Fig. 6.28) involve the data level and then require more of the hardware. Our current experimental setup (personal computer)⁷ allows us to reach a scale

⁷ These experiments were performed on a 2.3GHZ/4GB Intel Core i5 running Mac OS X 10.6.8 and MayBMS (a PostgreSQL 8.3.3 extension).

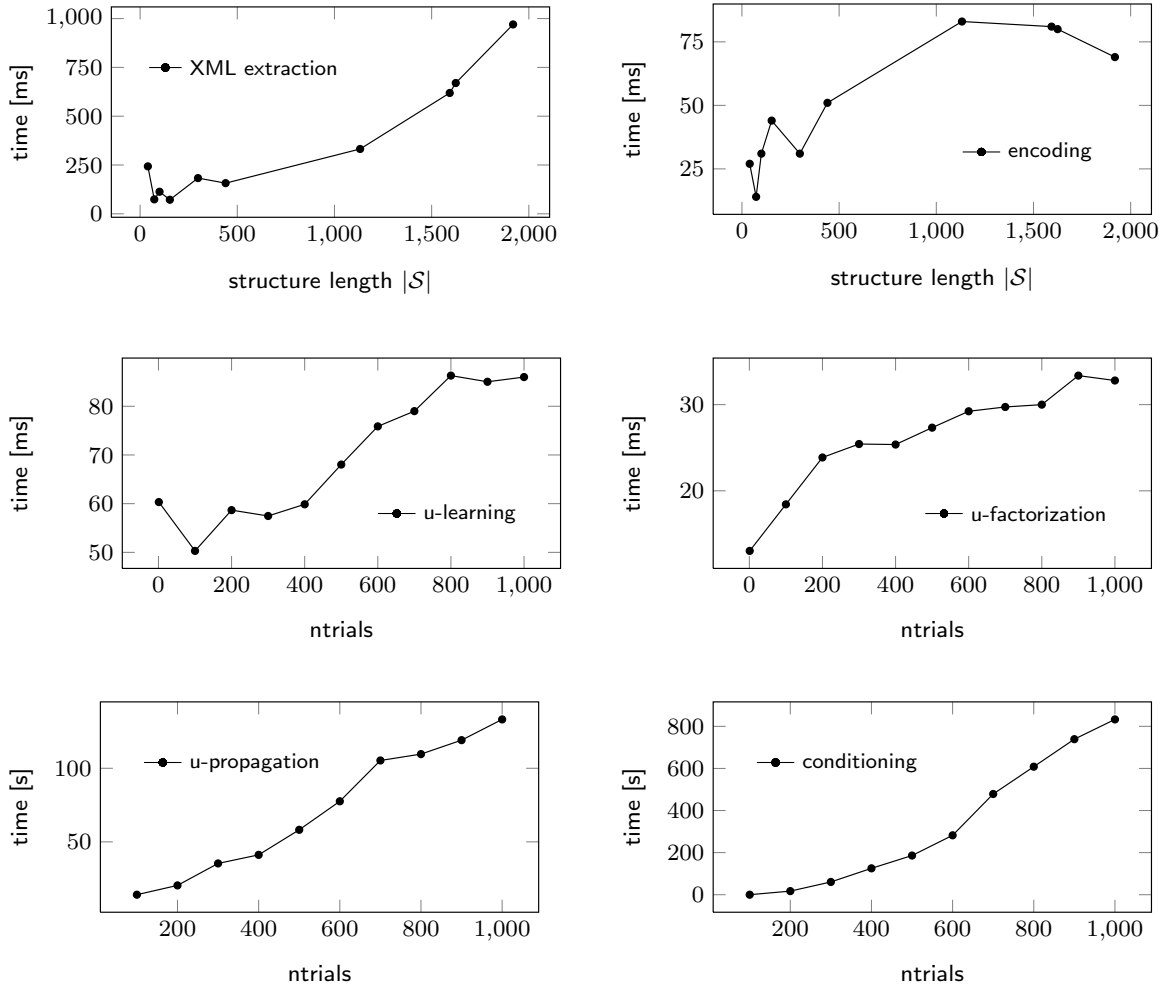


Figure 6.28. Performance behavior of Υ-DB on a Physiome testbed scenario.

in which the uncertain data being processed in synthesis ‘4U’ is sized up to 1 GB. For the two first graphs (*XML extraction* and *encoding*), we have collected the response time on the measure of interest over different structure lengths. Each one corresponds to a real Physiome hypothesis from the table of Fig. 6.29. The last hypothesis in that table, $v = 379$, is used for the tests of the four last graphs in Fig. 6.28, viz., *u-learning*, *u-factorization*, *u-propagation* and *conditioning*. We have set different number of trials (ntrials) over it, each one having 1 MB. The last test in each of such four graphs, with 1K trials, is processing 1 GB of uncertain data at once and then fits the machine’s main memory. We interpret the performance results shown in these graphs as follows for each measure of performance.

- *Extraction.* Some fluctuation may be due to practicalities of XML DOM access methods. The point of this performance study is to have practical

HYPOTHESIS	v	name	$ \mathcal{S} $	$ \mathcal{E} $
	186	Regulatory_Vessel	40	20
	89	Myo_Dyn_Resp_wFit	73	28
	60	Myogenic_Compliant_Vessel	100	38
	75	Baroreceptor_Lu_et_al_2001	153	74
	70	4-State_Sarcomere_Energetics	298	91
	120	Comp_four_gen_weibel_lung	440	186
	91	CardiopulmonaryMechanics	1132	412
	93	CardiopulmonMechGasBloodExch	1593	525
	153	HighlyIntegHuman	1624	538
	154	HighlyIntHuman_wIntervention	1919	634
	379	Baroreflex_SB_CT	171	74

Figure 6.29. Physiome hypotheses used in the experiments.

measures of the amount of time taken to process representative hypothesis structures. Note that even for structures of size $|\mathcal{S}| = 2K$ the amount of time required to extract a hypothesis is kept at subsecond order of magnitude (interactive response time) in a personal machine.

- *Encoding.* Some fluctuation is expected due to varying degrees of coupling between variables in the hypothesis structures. Note that, although $|\mathcal{S}|$ provides a very good measure of their size and complexity, the extent to which they are intricate should cause impact on the encoding procedure, which is any case kept $O(\sqrt{|\mathcal{E}||\mathcal{S}|})$. Again, the point here is to provide a notion of the amount of time required to encode representative hypotheses. For a scalability test on the encoding procedure, cf. Fig. 3.8.
- *U-intro.* The U-intro procedure is composed of *u-factor learning*, *u-factorization* and *u-propagation*. We observed in previous tests that it was dominated by the learning component, viz., the discovery of occasional fd's in the 'big' fact table. However, this is no longer the case once we implemented the workaround of keeping (in addition to the 'big' table) a table containing only the exogenous (input parameter) variables, as it has negligible size w.r.t. the data of the endogenous (predictive output) variables. Then u-learning became subsecond again and the U-intro procedure became dominated by *u-propagation*. In fact, the procedure of u-factorization, carried out once the fd's are discovered by u-learning, is also sub-second then

has negligible processing time w.r.t. u-propagation. The latter is the most expensive sub-procedure of the synthesis method.

- *Conditioning.* The conditioning procedure is run for a selected phenomenon. It is composed of four main parts. First, by operation `conf()` it performs a probabilistic inference sub-query on the proper predictive projection of the ‘big’ fact table of each hypothesis associated with the phenomenon. Second, it combines the results of each such sub-query through a union all query whose result set is a multi-hypothesis predictive table. Third, it loads the phenomenon observation sample data and the predictive data from the multi-hypothesis table into memory to apply Bayesian inference. Finally, the prior probability distribution of the predictive table is updated with the posterior and all the corresponding marginal probabilities are updated in their original U-relational tables. In our tests, this procedure is carried out over varying number of trials (`ntrials`). The total response times are shown in the last plot of Fig. 6.28.

This performance behavior is to be interpreted in the context of ETL in DW’s. Loading and setting up an Υ -DB has an overhead that shall be, though, much lower in high-performance machines. Such overhead is nonetheless justified for the use case of hypothesis management and analytics as opposed to simulation data management and exploratory analytics (cf. §2.6.2).

6.5. Discussion

We have verified that the hypothesis ratings/rankings shown in §6.2 coincide with the results (e.g., of model tuning) described in the Physiome model entries and their related publications. That validates the applicability of the Υ -DB methodology as a tool for data-driven analysis in such realistic scenarios.

The current practice in Computational Science for model evaluation and comparison in the presence of observational data is somewhat handcrafted: model agreement is assessed either qualitatively by referring to curve shapes in data plots or quantitatively by means of ad-hoc scripts. The Υ -DB methodology offers a tool to perform data-driven hypothesis analytics semi-automatically directly in the

database under the support of its querying capabilities. It has, therefore, potential to be a step towards higher standards of reproducibility and scalability.

Realistic assumptions. The core assumption of our framework is that the hypotheses are given in a formal specification which is encodable into a SEM that is complete (satisfies Defs. 1, 2). Also, as a semantic assumption which is standard in scientific modeling, we consider a one-to-one correspondence between real-world entities and variable/attribute symbols within a structure, and that all of them must appear in some of its equations/fd's. For most science use cases involving deterministic models (if not all), such assumptions are quite reasonable. It can be a topic of future work (cf. §7.3) to explore business use cases as well.

Hypothesis learning. The (user) method for hypothesis formation is irrelevant to our framework, as long as the resulting hypothesis is encodable into a SEM. So, a promising use case is to incorporate machine learning methods into our framework to scale up the formation/extraction of hypotheses and evaluate them under the querying capabilities of a p-DB. Consider, e.g., learning the equations, say, from Eureka [2].⁸

Qualitative hypotheses. The Υ -DB methodology is primarily motivated by computational science (usually involving differential equations). It is, however, applicable to qualitative deterministic models as well. Boolean Networks, e.g., consist in sets of functions $f(x_1, x_2, \dots, x_n)$, where f is a Boolean expression. For instance, Fig. 6.30 presents the system of Boolean equations of a tentative Boolean Network model for a plant hormone (Fig. 6.31) published in [78].⁹ The notation, e.g., **SphK***, is read (just like an ordinary differential equation), ‘the next state value of variable **SphK** is given by the state value of variable **ABA**. The parameters in this kind of model are the variable initial conditions.

⁸ <http://creativemachines.cornell.edu/Eureka>.

⁹ Cf. <http://atlas.bx.psu.edu/booleannet/booleannet.html>.

$$\mathcal{H} = \{ \begin{array}{ll} \text{SphK}^* & = \text{ABA}, \\ \text{S1P}^* & = \text{SphK}, \\ \text{GPA1}^* & = (\text{S1P or not GCR1}) \text{ and } \text{AGB1}, \\ \text{PLD}^* & = \text{GPA1}, \\ \text{PA}^* & = \text{PLD}, \\ \text{pHc}^* & = \text{ABA}, \\ \text{OST1}^* & = \text{ABA}, \\ \text{ROP2}^* & = \text{PA}, \\ \text{Atrboh}^* & = \text{pHc and OST1 and ROP2 and not ABI1}, \\ \text{ROS}^* & = \text{Atrboh} \\ \text{H+ATPase}^* & = \text{not ROS and not pHc and not Ca}^{2+}\text{c} \\ \text{ABI1}^* & = \text{pHc and not PA and not ROS}, \\ \text{RCN1}^* & = \text{ABA}, \\ \text{NIA12}^* & = \text{RCN1}, \\ \text{NOS}^* & = \text{Ca}^{2+}\text{c}, \\ \text{NO}^* & = \text{NIA12 and NOS}, \\ \text{GC}^* & = \text{NO}, \\ \text{ADPRc}^* & = \text{NO}, \\ \text{cADPR}^* & = \text{ADPRc}, \\ \text{cGMP}^* & = \text{GC}, \\ \text{PLC}^* & = \text{ABA and Ca}^{2+}\text{c}, \\ \text{InsP3}^* & = \text{PLC}, \\ \text{InsPK}^* & = \text{ABA}, \\ \text{InsP6}^* & = \text{InsPK}, \\ \text{CIS}^* & = (\text{cGMP and cADPR}) \text{ or } (\text{InsP3 and InsP6}), \\ \text{Ca}^{2+}\text{ATPase}^* & = \text{Ca}^{2+}\text{c}, \\ \text{Ca}^{2+}\text{c}^* & = (\text{CaIM or CIS}) \text{ and } (\text{not Ca}^{2+}\text{ATPase}), \\ \text{AnionEM}^* & = ((\text{Ca}^{2+}\text{c or pHc}) \text{ and not ABI1}) \text{ or } (\text{Ca}^{2+}\text{c and pHc}), \\ \text{Depolar}^* & = \text{KEV or AnionEM or (not H+ATPase) or (not KOUT) or Ca}^{2+}\text{c}, \\ \text{CaIM}^* & = (\text{ROS or not ERA1 or not ABH1}) \text{ and } (\text{not Depolar}), \\ \text{KOUT}^* & = (\text{pHc or not ROS or not NO}) \text{ and Depolar}, \\ \text{KAP}^* & = (\text{not pHc or not Ca}^{2+}\text{c}) \text{ and Depolar}, \\ \text{KEV}^* & = \text{Ca}^{2+}\text{c}, \\ \text{PEPC}^* & = \text{not ABA}, \\ \text{Malate}^* & = \text{PEPC and not ABA and not AnionEM}, \\ \text{RAC1}^* & = \text{not ABA and not ABI1}, \\ \text{Actin}^* & = \text{Ca}^{2+}\text{c or not RAC1}, \\ \text{Closure}^* & = (\text{KOUT or KAP}) \text{ and AnionEM and Actin and not Malate} \end{array} \}.$$

Figure 6.30. Example of Boolean Network hypothesis.

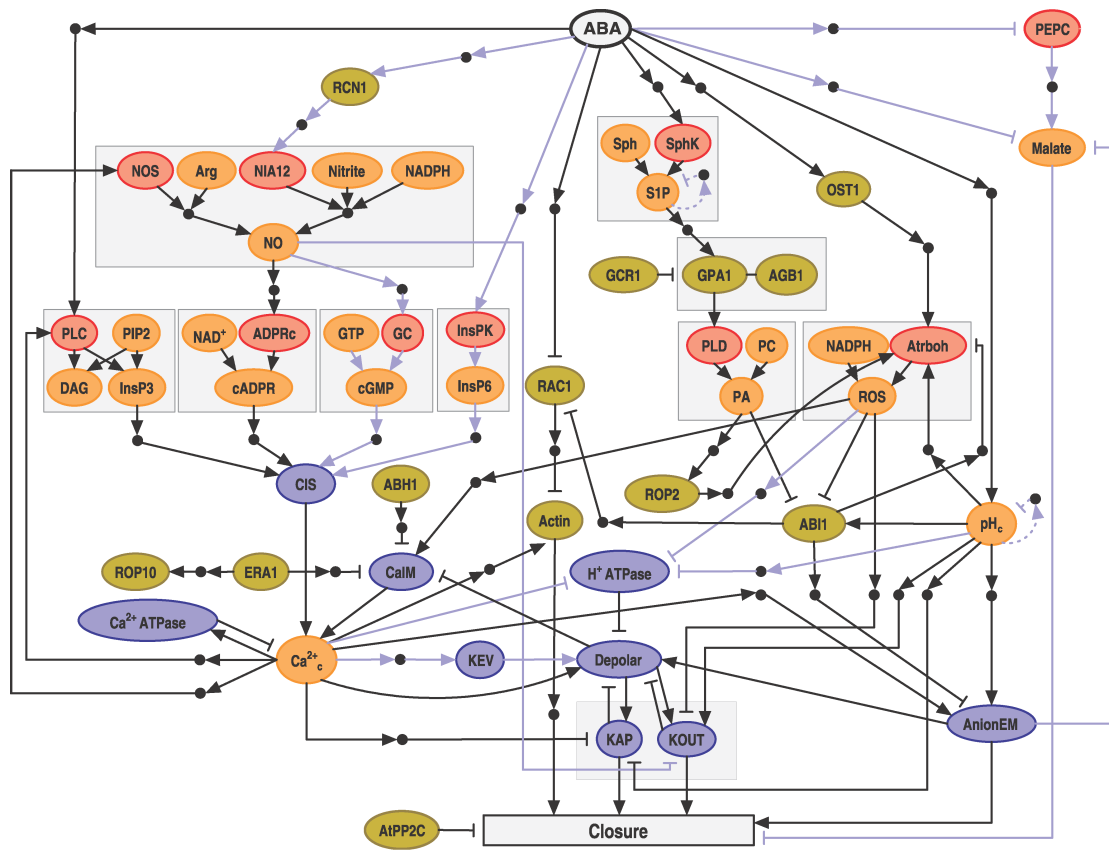


Figure 6.31. Example of Boolean Network model (source: [78]).

Several kinds of dynamical system can be modeled in this formalism. Applications have grown out of gene regulatory network to social network and stock market predictive analytics. Even if richer semantics is considered (e.g., fuzzy logics), our encoding method is applicable likewise, as long as the equations are still deterministic.

6.6. Conclusions

In this chapter we have demonstrated and discussed the applicability of the Υ -DB methodology. We have referred to real-world use case scenarios derived from the Physiome research project. We have shown in some detail the process of building an Υ -DB with representative models from Physiome's model repository. That qualitative assessment is followed by experiments that provide some concrete feel on the performance behavior of Υ -DB for models with up to 600+ mathematical variables.

Chapter 7

Conclusions

In this chapter we (§7.1) revisit the research questions addressed by this thesis, (§7.2) point out its significance and limitations, (§7.3) list open problems and topics for future work, and (§7.4) conclude with final considerations.

7.1. Revisiting the Research Questions

Let us now revisit the conceptual (RQ1-4) and technical (RQ5-9) research questions.

RQ1. *How to define and encode hypotheses ‘as data’? What are the sources of uncertainty that may be present and should be considered?*

In Chapter 2 we have provided core abstractions that compose the vision of hypotheses ‘as data,’ or the Υ -DB vision. The problem of hypothesis encoding has been defined and addressed further in Chapter 3. We have distinguished two main sources of uncertainty in our model of uncertainty for hypothesis management, viz., (i) *theoretical* uncertainty, as arising from competing hypotheses; and (ii) *empirical* uncertainty, as arising from the alternative trial datasets available for each hypothesis.

RQ2. *How does hypotheses ‘as data’ relate with observational data or, likewise, phenomena ‘as data’ from a database perspective?*

Also in Chapter 2, we have presented a conceptual framework in which we have defined hypotheses ‘as data’ and shown how it can be compared against phenomena ‘as data.’ In fact, hypothesis management is really

significant when it is possible to rate/rank hypotheses in the presence of (some partial piece of) evidence.

RQ3. *Does every piece of simulated data qualify as a scientific hypothesis? What is the difference between managing ‘simulation’ data from managing ‘hypotheses’ as data?*

Early in Table 1.1, we provided a comparison between simulation data management and hypothesis data management. Furthermore, the scientific research process is abstracted in Chapter 2 as a well-defined problem of data cleaning. Hypotheses are seen from an applied science point of view, and then are reduced into data such that a piece of simulation data is considered a hypothesis whenever it is assigned to explain some specific phenomenon.

RQ4. *Is there available a proper (machine-readable) data format we can use to automatically extract mathematically-expressed hypotheses from?*

We anticipated in Chapter 2 the adoption of the XML data model as the general data format for extracting hypothesis specifications from. In particular, since we deal here with mathematical hypotheses, we refer to **MathML** as a standard for hypothesis specification. Concretely, in Chapter 6 we present use case demonstration scenarios for which we have developed a specific wrapper, viz., for the extraction of hypotheses specified in **MML** (Mathematical Modeling Language).

RQ5. *Is there an algorithm to, given a SEM, efficiently extract its causal ordering? What are the computational properties of this problem?*

As shown in Chapter 3, Simon’s treatment of the problem of causal ordering given a SEM $\mathcal{S}(\mathcal{E}, \mathcal{V})$ is NP-Hard. In the same chapter, we have discussed this problem in detail and presented an effective, efficient algorithmic approach to the problem. The computational cost for the whole process of hypothesis encoding is bounded by $O(\sqrt{|\mathcal{S}|}|\mathcal{E}|)$. Experiments show that the approach performs well in practice for large hypotheses.

RQ6. *What is the connection between SEM's and fd's? Can we devise an encoding scheme to 'orient equations' and then effectively transform one into the other with guarantees? Once we do it, what design-theoretic properties have such a set of fd's?*

Also in Chapter 3, we have presented an algorithmic encoding scheme to transform a SEM into a set of fd's with guarantees in terms of preserving the hypothesis causal structure. Our study of this problem has revealed some interesting properties of the resulting fd sets, in particular, that they are always 'non-redundant' and, in comparison with arbitrary information systems, more precise and economical in the sense that, for any given attribute, there is exactly one fd with it in its rhs.

RQ7. *Is such fd set ready to be used for p-DB schema synthesis as an encoding of the hypothesis causal structure? If not, what kind of further processing we have to do? Can we perform it efficiently by reasoning directly on the fd's? How does it relate to the SEM's causal ordering?*

As we discuss in Chapter 4, the encoded fd set must be further processed to find the 'first causes' for each of its predictive variable. For addressing that, in Chapter 4 we have presented the concept of the folding of an fd set and an efficient algorithm to compute it. Also, we have shown the equivalence of such fd reasoning with causal ordering processing.

RQ8. *Is the uncertainty decomposition required for predictive analytics reducible to the structure level (fd processing), or do we need to process the simulated data to identify additional uncertainty factors? Finally, what properties are desirable for a p-DB schema targeted at hypothesis management? Are they ensured by this synthesis method?*

In Chapter 5 we have presented a conceptual framework to address synthesis for uncertainty '4U.' In particular, we have introduced the need to process, for each hypothesis, its trial datasets available, and presented an efficient algorithm to factorize and propagate the overall uncertainty present in a given hypothesis (as a competing explanation for a target

phenomenon). Then we have motivated BCNF as a notion of “good” design w.r.t. the factorized fd set based on the folding concept, and the lossless join property as required for the correctness of uncertainty decomposition. We have shown that the synthesized p-DB schema bears both properties.

RQ9. *Given all such a design-theoretic machinery to process hypotheses into (U-)relational DB's, what properties can we detect on the hypotheses back at the conceptual level? Do we have now technical means to speak of hypotheses that are “good” in terms of principles of the philosophy of science?*

Equipped with the design-theoretic machinery proposed in this thesis, we are able to, given a SEM, to automatically **(1)** extract its causal ordering, **(2)** detect its strongly coupled components and decide, for a given predictive projection, what are its associated u-factor projections (if any), and shall be able as well to **(3)** query the hypothesis ranking for a phenomenon of interest. All these are technical means to [15]: **(1')** extract the hypothesis ‘empirical content’ and ‘predictive power;’ **(2')** unravel its cohesiveness and how parsimonious it is in terms of the number of different claims or epistemological units carried within it, as well as its empirical grounding (‘first causes’); and finally, we shall be able to **(3')** appraise it in face of competing or alternative explanations.

7.2. Significance and Limitations

This thesis addresses the pressing call for large-scale, data-driven hypothesis management and analytics [35, 3, 10]. Some reasons that contribute for its significance are listed (cf. [10, 29, 28]).

- Structured deterministic hypotheses are now shown to be *encodable* as uncertain and probabilistic (U-relational) data based on p-DB principles;
 - * Study of the connection between SEM's and fd's, with contribution both to computational properties of the causal ordering problem, and to causal reasoning over fd's;

- * First *synthesis* method for the construction of p-DB's from some previous existing formal specification.
- Definition of a concrete *use case* of data-driven hypothesis management and analytics;
 - * New class of *applications* introduced for p-DB's;
 - * Settled the problem of Bayes' *conditioning* in p-DB's.

Now some limitations of the thesis are listed.

- The Bayesian inference is implemented at application level, yet not formulated as a principled technical solution within research in p-DB's.
- The encoding scheme to transform the mathematical system of a hypothesis into a set of fd's enabling the synthesis of the p-DB is applicable to structured *deterministic* models only, not stochastic ones.

7.3. Open Problems and Future Work

Open problems and topics of future work are listed (no particular order).

- (1) The design of a dedicated algebraic operation for Bayes' conditioning in p-WSA.
- (2) Investigation of other data dependency formalisms (e.g., multi-valued dependencies [20]), approximate fd's [68], conditional fd's [79]) to extend the scope of Υ -DB towards structured stochastic models.
- (3) Development of techniques for systematic hypothesis extraction as a well-defined problem of (web) information extraction;
- (4) Investigation of business use case scenarios for data-driven decision making on top of Υ -DB;
- (5) Definition of a machine learning use case scenario to industrialize hypothesis formation and assess Υ -DB's performance feasibility in such a scenario;

- (6) Development of automatic data sampling techniques to leverage the data definition of both hypotheses and phenomena in Υ -DB from a statistical point of view.

7.4. Final Considerations

In this thesis we have developed the vision of Υ -DB, which is essentially an abstraction of hypotheses as uncertain and probabilistic data. It comprises a design-theoretic methodology for the systematic construction and management of U-relational hypothesis DB's. It is meant to provide a principled approach to enable scientists and engineers to manage and evaluate (rate/rank) large-scale scientific hypotheses as theoretical data. We have addressed some core technical challenges over the Υ -DB vision in order to properly encode deterministic hypotheses as uncertain and probabilistic data.

As envisioned by Jim Gray [1], the scientific method has been shifting towards being operated as a data-driven discipline which is rapidly gaining ground [3]. In this thesis we have strived for proposing some core principles and techniques for enabling data-driven hypothesis management and analytics, opening a promising line of research in both probabilistic databases and simulation data management.

Bibliography

- [1] HEY, T.; TANSLEY, S.; TOLLE, K. **The fourth paradigm: Data-intensive scientific discovery**. Microsoft Research, 2009.
- [2] SCHMIDT, M.; LIPSON, H. Distilling free-form natural laws from experimental data. **Science**, Washington, v. 324, n. 5923, p. 81–5, 2009.
- [3] DHAR, V. Data science and prediction. **Communications of the ACM**, v. 56, n. 12, p. 64–73, 2013.
- [4] JAGADISH, H. V.; GEHRKE, J.; LABRINIDIS, A.; PAPAKONSTANTINO, Y.; PATEL, J. M.; RAMAKRISHNAN, R.; SHAHABI, C. Big data and its technical challenges. **Communications of the ACM**, v. 57, n. 7, p. 86–94, 2014.
- [5] BENIOFF, M. R.; LAZOWSKA, E. D. (Eds.). **Computational Science: Ensuring America’s competitiveness**. PITAC (US President’s Information Technology Advisory Committee), 2005.
- [6] PERLMAN, E.; BURNS, R.; LI, Y.; MENEVEAU, C. Data exploration of turbulence simulations using a database cluster. In: . c2007.
- [7] MARKRAM, H. The Blue Brain Project. **Nature Reviews Neuroscience**, v. 7, p. 153–60, 2006.
- [8] AILAMAKI, A.; KANTERE, V.; DASH, D. Managing scientific data. **Comm. ACM**, v. 53, n. 6, p. 68–78, 2010.
- [9] AHMAD, Y.; BURNS, R.; KAZHDAN, M.; MENEVEAU, C.; SZALAY, A.; TERZIS, A. Scientific data management at the Johns Hopkins Institute for

- Data Intensive Engineering and Science. **SIGMOD Record**, v. 39, n. 3, p. 18–23, 2010.
- [10] GONCALVES, B.; PORTO, F. Υ -DB: Managing scientific hypotheses as uncertain data. **PVLDB**, v. 7, n. 11, p. 959–62, 2014.
- [11] CUSHING, J. B. Beyond big data? **Computing in Science & Engineering**, v. 15, n. 5, p. 4–5, 2013.
- [12] WEINBERG, R. Point: Hypotheses first. **Nature**, London, v. 464, p. 678, 2010.
- [13] GOLUB, T. Counterpoint: Data first. **Nature**, London, v. 464, p. 679, 2010.
- [14] FRÉGNAC, Y.; LAURENT, G. Where is the brain in the Human Brain Project? **Nature**, London, v. 513, p. 27–9, 2014.
- [15] LOSEE, J. **A historical introduction to the philosophy of science**. Oxford Univ. Press, 4th ed., 2001.
- [16] BOLSTAD, W. M. **Introduction to bayesian statistics**. 2nd. ed. Wiley-Interscience, 2007.
- [17] SUCIU, D.; OLTEANU, D.; RÉ, C.; KOCH, C. **Probabilistic databases**. Morgan & Claypool Publishers, 2011.
- [18] KOCH, C. **MayBMS: A system for managing large uncertain and probabilistic databases**. In C. Aggarwal (ed.), *Managing and Mining Uncertain Data*, Chapter 6. Springer-Verlag, 2009.
- [19] DARWICHE, A. Bayesian networks. **Comm. ACM**, v. 53, n. 12, p. 80–90, 2010.
- [20] ULLMAN, J. **Principles of Databases and Knowledge-Base Systems**. Computer Science Press, 1988.
- [21] ABITEBOUL, S.; HULL, R.; VIANU, V. **Foundations of databases**. Addison-Wesley, 1995.

- [22] MAIER, D. **Theory of relational databases**. Computer Science Press, 1983.
- [23] BADIA, A.; LEMIRE, D. A call to arms: Revisiting database design. **SIGMOD Record**, v. 40, n. 3, p. 61–9, 2011.
- [24] SIMON, H. **Causal ordering and identifiability**. In Hood & Koopmans (eds.), *Studies in Econometric Methods*, Chapter 3, John Wiley & Sons, 1953.
- [25] PEARL, J. **Causality: Models, reasoning, and inference**. Cambridge Univ. Press, 2000.
- [26] CHANG, C.-H.; KAYED, M.; GIRGIS, M. R.; SHAALAN, K. A survey of web information extraction systems. **IEEE Transactions on Knowledge and Data Engineering**, v. 18, n. 10, p. 1411–28, 2006.
- [27] ANTOVA, L.; JANSEN, T.; KOCH, C.; OLTEANU, D. Fast and simple relational processing of uncertain data. In: **Proc. of IEEE ICDE**. 2008. p. 983–92.
- [28] GONCALVES, B.; SILVA, F. C.; PORTO, F. Υ -DB: A system for data-driven hypothesis management and analytics. Technical report, LNCC, 2015. (available at CoRR abs/1411.7419).
- [29] GONCALVES, B.; PORTO, F. Design-theoretic encoding of deterministic hypotheses as constraints and correlations in U-relational databases. Technical report, LNCC, 2015.
- [30] SIMON, H.; RESCHER, N. Cause and counterfactual. **Philosophy of Science**, v. 33, n. 4, p. 323–40, 1966.
- [31] HINES, M.; MORSE, T.; MIGLIORE, M.; CARNEVALE, N.; SHEPHERD, G. ModelDB: A database to support computational neuroscience. **J. Comput. Neurosci.**, v. 17, n. 1, p. 7–11, 2004.
- [32] CHELLIAH, V.; LAIBE, C.; Le Novère, N. BioModels Database: A repository of mathematical models of biological processes. **Method. Mol. Biol.**, , n. 1021, p. 189–99, 2013.

- [33] HUNTER, P. J.; BORG, T. K. Integration from proteins to organs: the Physiome Project. **Nat. Rev. Mol. Cell. Biol.**, v. 4, n. 3, p. 237–43, 2003.
- [34] BASSINGTHWAIGHTE, J. B. Strategies for the Physiome Project. **Ann. Biomed. Eng.**, v. 28, p. 1043–58, 2000.
- [35] HAAS, P.; MAGLIO, P.; SELINGER, P.; TAN, W. Data is dead... without what-if models. **PVLDB**, v. 4, n. 12, p. 1486–9, 2011.
- [36] RAHM, E.; Hai Do, H. Data cleaning: Problems and current approaches. **IEEE Data Engineering Bulletin**, v. 23, n. 4, p. 3–13, 2001.
- [37] BESKALES, G.; SOLIMAN, M. A.; ILYAS, I. F.; BEN-DAVID, S. Modeling and querying possible repairs in duplicate detection. **PVLDB**, v. 2, n. 1, p. 598–609, 2009.
- [38] PORTO, F.; SPACAPPIETRA, S. **Data model for scientific models and hypotheses**. In R. Kaschek, L. Delcambre (ed.), *The evolution of Conceptual Modeling*, p. 285–305, Springer-Verlag LNCS vol. 6520, January 2011.
- [39] STOUGIANNIS, A.; TAUHEED, F.; PAVLOVIC, M.; HEINIS, T.; AILAMAKI, A. Data-driven Neuroscience: Enabling breakthroughs via innovative data management. In: . c2013.
- [40] KANOV, K.; PERLMAN, E. A.; BURNS, R. C.; AHMAD, Y.; SZALAY, A. S. I/O streaming evaluation of batch queries for data-intensive computational turbulence. In: . c2011. p. 1–10.
- [41] IDREOS, S.; ALAGIANNIS, I.; JOHNSON, R.; AILAMAKI, A. Here are my data files. here are my queries. where are my results? In: **In Proc. of the biennial Conf. on Innovative Data Systems Research (CIDR)**. 2011.
- [42] ALAGIANNIS, I.; BOROVICA, R.; BRANCO, M.; IDREOS, S.; AILAMAKI, A. NoDB: Efficient query execution on raw data files. In: **Proc. of ACM SIGMOD**. 2012. p. 241–52.
- [43] SOLDATOVA, L.; RZHETSKY, A. Representation of research hypotheses. **J. Biomed. Sem.**, v. 2, n. S2, 2011.

- [44] KING, R. D. et al. The automation of science. **Science**, Washington, v. 324, n. 5923, p. 85–9, 2009.
- [45] RACUNAS, S. et al. HyBrow: a prototype system for computer-aided hypothesis evaluation. **Bioinformatics**, v. 20, n. 1, p. 257–64, 2004.
- [46] GAO, Y. et al. SWAN: A distributed knowledge infrastructure for alzheimer disease research. **J. Web Semantics**, v. 4, n. 3, p. 222–8, 2006.
- [47] de WAARD, A. et al. Hypotheses, evidence and relationships: The Hyper approach for representing scientific knowledge claims. In: **ISWC Proc. of Workshop on Semantic Web Applications in Scientific Discourse**. 2009.
- [48] DASH, D.; DRUZDZEL, M. J. A note on the correctness of the causal ordering algorithm. **Artif. Intell.**, v. 172, n. 15, p. 1800–8, 2008.
- [49] NAYAK, P. P. Causal approximations. **Artificial Intelligence**, Amsterdam, v. 70, n. 1-2, p. 277–334, 1994.
- [50] DRUZDZEL, M. J.; SIMON, H. A. Causality in Bayesian belief networks. In: **Proc. of Int. Conf. on Uncertainty in Artificial Intelligence**. 1993. p. 3–11.
- [51] MADEIRA, S. C.; OLIVEIRA, A. L. Biclustering algorithms for biological data analysis: A survey. **IEEE Transactions on Computational Biology and Bioinformatics**, v. 1, n. 1, p. 24–45, 2004.
- [52] DHILLON, I. S.; MALLELA, S.; MODHA, D. S. Information-theoretic co-clustering. In: **Proc. of ACM SIGKDD**. 2003. p. 89–98.
- [53] EVEN, S. **Graph algorithms**. 2nd. ed. Cambridge Univ. Press, 2011.
- [54] UNO, T. An efficient algorithm for solving pseudo clique enumeration problem. **Algorithmica**, v. 56, n. 1, p. 3–16, 2010.
- [55] GAREY, M. R.; JOHNSON, D. S. **Computers and intractability: A guide to the theory of NP-completeness**. 1st. ed., Series of Books in the Mathematical Sciences. W. H. Freeman, 1979.

- [56] JOHNSON, D. S. The NP-completeness column: an ongoing guide. **J. Algorithms**, v. 8, n. 3, p. 438–88, 1987.
- [57] HOPCROFT, J. E.; KARP, R. M. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. **SIAM Journal on Computing**, v. 2, n. 4, p. 225–31, 1973.
- [58] CORMEN, T. H.; LEISERSON, C. E.; RIVEST, R. L.; STEIN, C. **Introduction to algorithms**. 3rd. ed. The MIT Press, 2009.
- [59] SIMON, H. On the definition of the causal relation. **The Journal of Philosophy**, v. 49, p. 517–28, 1952.
- [60] NAYAK, P. P. **Automated modelling of physical systems**. Springer-Verlag, 1996.
- [61] SERRANO, D.; GOSSARD, D. C. Constraint management in conceptual design. In: **Knowledge Based Expert Systems in Engineering: Planning and Design**. Computational Mechanics Publications, 1987. p. 211–24.
- [62] BERNSTEIN, P. Synthesizing third normal form relations from functional dependencies. **ACM Trans. on Database Systems**, v. 1, n. 4, p. 277–98, 1976.
- [63] BEERI, C.; BERNSTEIN, P. Computational problems related to the design of normal form relational schemas. **ACM Trans. on Database Systems**, v. 4, n. 1, p. 30–59, 1979.
- [64] MELIOU, A. et al. Causality in databases. **IEEE Data Eng. Bull.**, v. 33, n. 3, p. 59–67, 2010.
- [65] MELIOU, A.; GATTERBAUER, W.; MOORE, K. F.; SUCIU, D. The complexity of causality and responsibility for query answers and non-answers. **PVLDB**, v. 4, n. 1, p. 34–45, 2010b.
- [66] KANAGAL, B.; LI, J.; DESHPANDE, A. Sensitivity analysis and explanations for robust query evaluation in probabilistic databases. In: **Proc. of ACM SIGMOD**. 2011. p. 841–52.

- [67] MELIOU, A.; GATTERBAUER, W.; SUCIU, D. Reverse data management. **PVLDB**, v. 4, n. 12, p. 1490–3, 2011.
- [68] HUHTALA, Y. et al. TANE: An efficient algorithm for discovering functional and approximate dependencies. **Computer Journal**, v. 42, n. 2, p. 100–11, 1999.
- [69] DARWICHE, A. **Modeling and reasoning with Bayesian Networks**. Cambridge University Press, 2009.
- [70] TOM S. VERMA, J. P. Causal Networks: Semantics and expressiveness. In: **Proc. of the 4th Conf. on Uncertainty in Artificial Intelligence (UAI'88)**. North-Holland Publishing Co., 1988. p. 69–78.
- [71] Das Sarma, A.; ULLMAN, J.; WIDOM, J. Schema design for uncertain databases. In: **Proc. of AMW**. 2007.
- [72] FAGIN, R. The decomposition versus synthetic approach to relational database design. In: **Proc. of VLDB**. 1977. v. 3, p. 441–6.
- [73] JAGADISH, H. V. et al. Making database systems usable. In: **SIGMOD**. 2007. p. 13–24.
- [74] QIAN, L.; LEFEVRE, K.; JAGADISH, H. V. CRIUS: User-friendly database design. **PVLDB**, v. 4, n. 2, p. 81–92, 2010.
- [75] KOCH, C.; OLTEANU, D. Conditioning probabilistic databases. **PVLDB**, v. 1, n. 1, p. 313–25, 2008.
- [76] BUGENHAGEN, S. M.; COWLEY, A. W. J.; BEARD, D. A. Identifying physiological origins of baroreflex dysfunction in salt-sensitive hypertension in the Dahl SS rat. **Physiological Genomics**, v. 42, p. 23–41, 2010.
- [77] ELTON, C.; NICHOLSON, M. The ten-year cycle in numbers of the lynx in Canada. **Journal of Animal Ecology**, v. 11, n. 2, p. 215–44, 1942.
- [78] LI, S.; ASSMANN, S. M.; ALBERT, R. Predicting essential components of signal transduction networks: A dynamic model of guard cell abscisic acid signaling. **PLOS Biology**, v. 4, n. 10, p. e312, 2006.

-
- [79] FAN, W.; GEERTS, F.; LI, J.; XIONG, M. Discovering conditional functional dependencies. **IEEE Transactions on Knowledge and Data Engineering**, v. 23, n. 5, p. 683–98, 2011.

Appendix A

Detailed Proofs

A.1. Proofs of Hypothesis Encoding

A.1.1 Proof of Theorem 1

“Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure. Then the extraction of its causal ordering by Simon’s COA(\mathcal{S}) is intractable (NP-Hard).”

Proof 17 We show that, at each recursive step of COA, to find all non-trivial minimal subsets (i.e., $|\mathcal{E}'| \geq 2$) translates into an optimization problem associated with the decision problem BPBP, which we know by Lemma 1 to be NP-Complete.

First, recall (Def. 2) that a structure $\mathcal{S}(\mathcal{E}, \mathcal{V})$ is complete if $|\mathcal{E}| = |\mathcal{V}|$; e.g., for the structure given in Fig. 3.5 (left), note (Def. 9) the minimal structure $\mathcal{S}'(\mathcal{E}', \mathcal{V}')$, where $\mathcal{E}' = \{f_1, f_2, f_3\}$. For non-trivial minimal structures, i.e., when $|\mathcal{E}'| = K \geq 2$, it is easy to see that its corresponding bipartite graph $G = (V'_1 \cup V'_2, E')$, where $V_1 \mapsto \mathcal{E}$, $V_2 \mapsto \mathcal{V}$ and $E \mapsto \mathcal{S}$ must have number of edges $|E'| \geq 2K$ and, for all its vertices $u \in V'_1 \cup V'_2$, u must have $\deg(u) \geq 2$, i.e., G is a pseudo-biclique in accordance with Def. 6. That intuition is elaborated as follows.

The point is that, no matter how big is such structure \mathcal{S}' , its equations $f \in \mathcal{E}'$ are such that $|\text{Vars}(f)| \geq 2$ (as \mathcal{S}' is non-trivial) and its variables can be grouped in local patterns from the sparsest kind to the densest. To construct an instance of the sparsest case, let \mathcal{S}' be built by setting a first equation where its entry in the structure matrix $A_{\mathcal{S}}$ has form $(1, 1, 0^+)$ and then, for the next $|\mathcal{E}'| - 2$ equations, shift such pair of 1’s one position right w.r.t. the previous one. Then complete it with a last equation whose form is form $(1, 0^+, 1)$. That is, the structure is built

with unique pairs of 1's spread all over the structure. Then, deciding whether there is a minimal structure of size $|\mathcal{E}'| = K$ corresponds exactly to BPBP. It is a special case (BBP), when such minimal structure is the densest possible, i.e., when A_S has only 1's and then its corresponding bipartite graph is a K -balanced biclique with $|E'| = K^2$, and $\deg(u) = K$ for all vertices $u \in V'_1 \cup V'_2$. For instance, see the minimal structure with $\mathcal{E}' = \{f_4, f_5\}$ found at the second recursive step of COA in Fig. 3.2. \square

A.1.2 Proof of Proposition 1

“Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a structure, and $\varphi_1 : \mathcal{E} \rightarrow \mathcal{V}$ and $\varphi_2 : \mathcal{E} \rightarrow \mathcal{V}$ be any two total causal mappings over \mathcal{S} . Then $C_1^+ = C_2^+$.”

Proof 18 The proof is based on an argument from Nayak [49], which we present here arguably much clearer. Intuitively, it shows that if φ_1 and φ_2 differ on the variable an equation f is mapped to, then such variables, viz., $\varphi_1(f)$ and $\varphi_2(f)$, must be causally dependent on each other (strongly coupled).

To show $C_1^+ = C_2^+$ reduces to $C_1^+ \subseteq C_2^+$ and $C_2^+ \subseteq C_1^+$. We show the first containment, with the second being understood as following by symmetry. Closure operators are extensive, $X \subseteq cl(X)$, and idempotent, $cl(cl(X)) = cl(X)$. That is, if we have $C_1 \subseteq C_2^+$, then we shall have $C_1^+ \subseteq (C_2^+)^+$ and, by idempotence, $C_1^+ \subseteq C_2^+$.

Then it suffices to show that $C_1 \subseteq C_2^+$, i.e., for any $(x', x) \in C_1$, we must show that $(x', x) \in C_2^+$ as well. Observe by Def. 5 that both φ_1 and φ_2 are bijections, then, invertible functions. If $\varphi_1^{-1}(x) = \varphi_2^{-1}(x)$, then we have $(x', x) \in C_2$ and thus, trivially, $(x', x) \in C_2^+$. Else, φ_1 and φ_2 disagree in which equations they map onto x . But we show next, in any case, that we shall have $(x', x) \in C_2^+$.

Take all equations $g \in \mathcal{E}' \subseteq \mathcal{E}$ such that $\varphi_1(g) \neq \varphi_2(g)$, and let $n \leq |\mathcal{E}|$ be the number of such ‘disagreed’ equations. Now, let $f \in \mathcal{E}'$ be such that its mapped variable is $x = \varphi_1(f)$. Construct a sequence of length $2n$ such that, $s_0 = \varphi_1(f) = x$ and, for $1 \leq i \leq 2n$, element s_i is defined $s_i = \varphi_2(\varphi_1^{-1}(s_{i-1}))$. That is, we are defining the sequence such that, for each equation $g \in \mathcal{E}'$, its disagreed mappings $\varphi_1(g) = x_a$ and $\varphi_2(g) = x_b$ are such that $\varphi_1(g)$ is immediately followed by $\varphi_2(g)$. As $x_a, x_b \in Vars(g)$, we have $(x_a, x_b) \in C_2$ and, symmetrically, $(x_b, x_a) \in C_1$.

The sequence is of form $s = \langle \underbrace{x, x_f, \dots, x_a}_f, \underbrace{x_b, \dots, x_{2n-1}}_g, \underbrace{x_{2n}}_h \rangle$.

Since x must be in the codomain of φ_2 , we must have a repetition of x at some point $2 \leq k \leq 2n$ in the sequence index, with $s_k = x$ and $s_{k-1} = x''$ such that $(x'', x) \in C_2$. If $x'' = x'$, then $(x', x) \in C_2$ and obviously $(x', x) \in C_2^+$. Else, note that x_f must also be in the codomain of φ_1 , while x'' in the codomain of φ_2 . Let ℓ be the point in the sequence, $3 \leq \ell \leq 2n-1$, at which $s_\ell = x_f = x_a$ and $s_{\ell+1} = x_b$ for some x_b such that $(x_f, x_b) \in C_2$. It is easy to see that, either we have $x_b = x''$ or $x_b \neq x''$ but $(x_b, x'') \in C_2^+$. Thus, by transitivity on such a causal chain, we must have $(x_f, x'') \in C_2^+$ and eventually $(x_f, x) \in C_2^+$. Finally, since $x' \in Vars(f)$ and $\varphi_2(f) = x_f$, we have $(x', x_f) \in C_2$ and, by transitivity, $(x', x) \in C_2^+$. \square

A.1.3 Proof of Theorem 2.

“Let Σ be an fd set defined $\Sigma \triangleq \text{h-encode}(\mathcal{S})$ for some complete structure \mathcal{S} . Then Σ is non-redundant and singleton-rhs but may not be left-reduced (then may not be canonical).”

Proof 19 We will show that properties (a-b) of Def. 9 hold for Σ produced by (Alg. 3) **h-encode**, but property (c) may not hold.

At initialization, the algorithm sets $\Sigma = \emptyset$ and then inserts an fd $\langle X, A \rangle \in \Sigma$ for each $\langle f, x \rangle \in \varphi_t$ scanned, where $x \mapsto A$ and $X \cap A = \emptyset$. At termination, for all fd's in Σ we obviously have $|A| = 1$ then property (a) holds. Also, note that $\varphi: S \rightarrow Vars(S)$ is, by Def. 5, a bijection.

Now, for property (b) not to hold there must be some fd $\langle X, A \rangle \in \Sigma$ that is redundant and then can be found in the closure of $\Gamma = \Sigma \setminus \langle X, A \rangle$. By Lemma 4 (below), that can be the case only if $A \subseteq X$ or there is $\langle Y, A \rangle \in \Gamma$ for some Y . But from $X \cap A = \emptyset$, we have $A \not\subseteq X$; and from φ being a bijection it follows that there can be no such fd in Γ . Thus it must be the case that Σ is non-redundant, i.e., property (b) holds.

Finally, property (c) does not hold if there can be some fd $\langle X, A \rangle \in \Sigma$ with $Y \subset X$ such that $\Gamma = \Sigma \setminus \langle X, A \rangle \cup \langle Y, A \rangle$ has the same closure as Σ . That is, if we may find $\langle Y, A \rangle \in \Sigma^+$. Now, pick structure S whose (3×3) matrix A_s has rows $(1, 0, 0), (1, 1, 0), (1, 1, 1)$ as an instance. Alg. 3 encodes it into $\Sigma = \{\phi \rightarrow x_1, x_1 v \rightarrow$

$x_2, x_1 x_2 v \rightarrow x_3\}$. Let $Y = \{x_1, v\}$, and $B = \{x_2\}$. Note that $x_1 v \rightarrow x_2 \in \Sigma$ can be written as $\langle Y, B \rangle \in \Sigma$, and $x_1 x_2 v \rightarrow x_3 \in \Sigma$ as $\langle YB, A \rangle \in \Sigma$. Now observe that $\langle Y, A \rangle \in \Sigma^+$ can be derived by R5 over $\langle Y, B \rangle, \langle YB, A \rangle \in \Sigma$, which is sufficient to show that property (c) may not hold. That is, B is “extraneous” in $\langle YB, A \rangle \in \Sigma$ and can be removed from its lhs without loss of information to Σ . \square

Lemma 4 Let Σ be a (Def. 9-a) singleton-rhs fd set on attributes U . Then $X \rightarrow A$ can only be in Σ^+ , where $XA \subseteq U$, if $A \subseteq X$ or there is non-trivial $\langle Y, A \rangle \in \Sigma$ for some $Y \subset U$.

Proof 20 By Lemma 5 (below), we know that $X \rightarrow A \in \Sigma^+$ iff $A \subseteq X^+$. We need to prove that if $A \not\subseteq X$ and there is no $Y \rightarrow A$ in singleton-rhs Σ , then $A \not\subseteq X^+$. But this is equivalent to show that (Alg. 4) **XClosure** gives only correct answers for X^+ w.r.t. Σ , which is known (cf. theorem from Ullman [20, p. 389]). Note that **XClosure**(Σ, X) inserts A in X^+ only if $A \subseteq X$ or there is some fd $\langle Y, A \rangle \in \Sigma$. \square

Lemma 5 Let Σ be an fd set. An fd $X \rightarrow Y$ is in Σ^+ iff $Y \subseteq X^+$, where X^+ is the attribute closure of X w.r.t. Σ .

Proof 21 This is from Ullman [20, p. 386]. Let $Y = A_1 \dots A_n$ and suppose $Y \subseteq X^+$. Then for each A_i , we have $A_i \in X^+$ and, by the definition of X^+ , we must have $\langle X, A_i \rangle \in \Sigma^+$. Then it follows by (R4) union that $X \rightarrow Y$ is in Σ^+ as well. Conversely, suppose $\langle X, Y \rangle \in \Sigma^+$. Then, by (R3) decomposition we have $\langle X, A_i \rangle \in \Sigma^+$ for each $A_i \in Y$. \square

A.2. Proofs of Causal Reasoning

A.2.1 Proof of Lemma 2

“Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, φ a total causal mapping over \mathcal{S} and Σ an fd set encoded through φ given \mathcal{S} . If $\langle X, A \rangle \in \Sigma$, then A^{q} , the attribute folding of A (w.r.t. Σ) exists and is unique.”

Proof 22 The existence of A^{q} is ensured by the degenerate case where $X = A^{\text{q}}$, as $X \rightarrow A$ is itself in Σ^{\triangleright} by an empty application of R5. If $X \rightarrow A$ is in fact folded w.r.t. Σ , then the folding of A exists. Else, it is not folded yet $X \rightarrow A$ is non-trivial

because by Theorem 2 Σ is non-redundant. Then, by Def. 11 there must be some $Y \subseteq U$ with $Y \not\supseteq X$ such that $Y \rightarrow X$ is in Σ^+ and $X \not\rightarrow Y$. By Def. 10, there is a finite application of R5 over fd's in Σ to derive $Y \xrightarrow{\supset} X$. Then by R2 \sim R5 over $X \rightarrow A$, we have $Y \rightarrow A$. Although there may be many such (intermediate) attribute sets $Y \subset U$ along the transitive chaining satisfying the conditions above, we claim there is at least one that is a folding of A . Suppose not. Then, for all such $Y \subset U$, there is some $Y' \subset U$ with $Y' \not\supseteq Y$ such that $Y' \rightarrow Y$ and $Y \not\rightarrow Y'$, leading to an infinite regress. Nonetheless, in so far as cycles are ruled out by force of Def. 11, then Σ^+ must have an infinite number of fd's. But Σ^+ is finite, viz., bounded by $2^{2|U|}$ (cf. [21, p. 165]). \nexists . Therefore the folding of A must exist.

Moreover, observe that Σ is encoded through φ , which is by Def. 5 a bijection. Then we have $\langle X, A \rangle \in \Sigma$ for exactly one attribute set X . Then, as a straightforward follow-up of the rationale that led us to infer the folding existence, note that there must be a single chaining $Y^n \xrightarrow{\supset} \dots \xrightarrow{\supset} Y^1 \xrightarrow{\supset} Y^0 \xrightarrow{\supset} X \xrightarrow{\supset} A$. Again, as cycles are ruled out by force of Def. 11 and Σ^+ is finite, then the folding of A is unique. \square

A.2.2 Proof of Theorem 3

“Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, and Σ an fd set encoded given \mathcal{S} . Now, let $\langle X, A \rangle \in \Sigma$. Then $\text{AFolding}(\Sigma, A)$ correctly computes A^{\supset} , the attribute folding of A (w.r.t. Σ) in time $O(|\mathcal{S}|^2)$.”

Proof 23 For the proof roadmap, note that AFolding is monotone (size of A^* can only increase) and terminates precisely when $A^{(i+1)} = A^{(i)}$, where $A^{(i)}$ denotes the attributes in A^* at step i of the outer loop. The folding A^{\supset} of A at step i is $A^{(i)} \setminus \Lambda^{(i)}$. We shall prove by induction, given attribute A from fd $X \rightarrow A$ in Σ , that $A^* \setminus \Lambda$ returned by $\text{AFolding}(\Sigma, A)$ is the unique attribute folding A^{\supset} of A .

(Base case). By Theorem 2, Σ is non-redundant with (then) non-trivial $\langle X, A \rangle \in \Sigma$ for exactly one attribute set X , the algorithm always reaches step $i = 1$, which is our base case. Then X is placed in $A^{(1)}$ and A in $\Lambda^{(1)}$, and we have $A^{(1)} = XA$ and $\Lambda^{(1)} = A$. Therefore, $A^{(1)} \setminus \Lambda^{(1)} = X$, and in fact we have $\langle X, A \rangle \in \Sigma^{\supset}$ by an empty application of R5. For it to be specifically in $\Sigma^{\supset} \subset \Sigma^{\supset}$, it must be folded w.r.t. set Δ of consumed fd's at this step, viz., $\Delta^{(1)} = \{X \rightarrow A\}$. In fact, as the only fd in $\Delta^{(1)}$, by Def. 11 it must be folded w.r.t. $\Delta^{(1)}$, and we

have $A^{\triangleright} = X$ at step $i=1$.

(Induction). Now, let $i = k$, for $k > 1$, and assume that $\langle A^{(k)} \setminus \Lambda^{(k)}, A \rangle \in \Sigma^{\triangleright} \subset \Sigma^{\triangleright}$ with $A^{(k)} \neq \Lambda^{(k)}$. By Lemma 2 we know that $A^{\triangleright} = A^{(k)} \setminus \Lambda^{(k)}$ is the unique folding of A at step $i = k$. For the inductive step, suppose Y is placed in $A^{(k+1)}$ and B in $\Lambda^{(k+1)}$ because $\langle Y, B \rangle \in \Sigma \setminus \Delta^{(k)}$ and $B \in A^{(k)}$.

Since $B \in A^{(k)}$ and $B \notin \Lambda^{(k)}$ (it is yet just be consumed into $\Lambda^{(k+1)}$), we can write $(A^{(k)} \setminus \Lambda^{(k)}) = ZB$ for some $Z \neq B$, where $\underbrace{(A^{(k)} \setminus \Lambda^{(k)})}_{ZB} \rightarrow A$ is assumed in Σ^{\triangleright} . Now, with the application of R5 consuming $Y \rightarrow B$ we have $\underbrace{(A^{(k)}YB \setminus \Lambda^{(k)}B)}_{ZS} \rightarrow A$, where $S = Y \setminus \Lambda^{(k)}$. We claim that $ZS \rightarrow A$ is folded w.r.t. $\Delta^{(k+1)}$.

Suppose not. Then by Def. 11 there must be some $W \not\subseteq ZS$ such that $W \rightarrow ZS$ is in $(\Delta^{(k+1)})^+$ but $ZS \not\rightarrow W$. Since $ZS \neq \emptyset$, there must be some $C \in ZS$, i.e., $C \notin \Lambda^{(k+1)}$. Note that, as $W \rightarrow ZS$ is in $(\Delta^{(k+1)})^+$, then by (R3) decomposition we have $W \rightarrow C$ in $(\Delta^{(k+1)})^+$ as well. But by Lemma 4 that can only be the case if there is some $\langle T, C \rangle \in \Delta^{(k+1)}$, which means C has been already consumed into $\Lambda^{(k+1)}$, though $C \notin \Lambda^{(k+1)}$. \downarrow .

Finally, as for the time bound, note that in worst case, exactly one fd $Y \rightarrow B$ is consumed from Σ into Δ for each step of the outer loop, where $|\Sigma| = |\mathcal{E}|$. That is, let $n = |\mathcal{E}|$, then n is decreased stepwise in arithmetic progression such that $n + (n-1) + \dots + 1 = n(n-1)/2$ scans are required overall, i.e., $O(n^2)$. Note also, however, that B may be the only symbol read at each such fd scan but in worst case at most $|U| = |\mathcal{V}|$ symbols are read. Thus our measure n should be actually overestimated $n = |\mathcal{E}| |\mathcal{V}| = |\mathcal{S}|$. Therefore Alg. 6 is bounded by $O(|\mathcal{S}|^2)$. \square

A.2.3 Proof of Corollary 2

*“Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, and Σ an fd set encoded given \mathcal{S} . Then algorithm **folding**(Σ) correctly computes Σ^{\triangleright} , the folding of Σ in time that is $f(|\mathcal{S}|) \Theta(|\mathcal{E}|)$, where $f(|\mathcal{S}|)$ is the time complexity of (Alg. 6) **AFolding**.”*

Proof 24 By Theorem 3, we know that sub-procedure (Alg. 6) **AFolding** is correct and terminates. Then (Alg. 5) **folding** necessarily inserts in Σ^{\triangleright} (initialized empty) exactly one fd $Z \xrightarrow{\triangleright} A$ for each fd $X \rightarrow A$ in Σ scanned. Thus, at termination we have $|\Sigma^{\triangleright}| = |\Sigma|$. Again, as **AFolding** is correct, we know Z is the unique folding of

A. Therefore it must be the case that Alg. 5 is correct. Finally, for the time bound, the algorithm iterates over each fd in Σ without having to read its symbols, and at each such step **AFolding** takes time that is $f(n)$. Thus **folding** takes $f(|\mathcal{S}|) \Theta(|\mathcal{E}|)$. But we know from Theorem 3 and Remark 4 that $f(|\mathcal{S}|) \in O(|\mathcal{S}|)$, then it takes $O(|\mathcal{S}| |\mathcal{E}|)$. \square

A.2.4 Proof of Proposition 3

“Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, φ a total causal mapping over \mathcal{S} and Σ an fd set encoded through φ given \mathcal{S} . Let Σ^{q} be the folding of Σ , then Σ^{q} is parsimonious.”

Proof 25 By Lemma 2 we know that, for each fd $\langle X, A \rangle \in \Sigma$, the attribute folding Z of A such that $Z \xrightarrow{\text{q}} A$ exists and is unique. That is, for no $Y \neq Z$ we have $Y \xrightarrow{\text{q}} A$. Thus $\Sigma^{\text{q}} \triangleq \text{folding}(\Sigma)$ automatically satisfies Def. 13, as long as we show it is canonical (cf. Def. 9).

Moreover, by Theorem 2 we know that Σ is both non-redundant and singleton-rhs. Now, consider by Lemma 2 that **AFolding** builds a bijection mapping each $\langle X, A \rangle \in \Sigma$ to exactly one $\langle Z, A \rangle \in \Sigma^{\text{q}}$ such that $Z \xrightarrow{\text{q}} A$. Since Σ is singleton-rhs, it is obvious that Σ^{q} is as well and covers all attributes in the rhs of fd's in Σ . Also, the bijection implies $|\Sigma^{\text{q}}| = |\Sigma|$. Since Σ is non-redundant and has exactly one fd with each attribute A in its rhs, then by Lemma 4 so is Σ^{q} .

Finally we will show that unlike Σ , its folding Σ^{q} must be left-reduced. Suppose not. Then for some fd $Z \rightarrow A$ in Σ^{q} there is $S \subset Z$ such that non-trivial $S \rightarrow A$ is in $(\Sigma^{\text{q}})^+$. Since $Z \rightarrow A$ is the only fd in Σ^{q} with A in its rhs and $S \rightarrow A$ is non-trivial, we must have $S \xrightarrow{\Delta} Z \xrightarrow{\Delta} A$.

Now, suppose $S \rightarrow A$ is not folded. Then there is $W \not\subseteq S$ such that $W \rightarrow S$ is in $(\Sigma^{\text{q}})^+$ but $S \not\rightarrow W$. Note that $W \neq Z$, as $W \not\subseteq S$. Also, $W \rightarrow S$ and $S \rightarrow Z$ implies $W \rightarrow Z$ by (R5) transitivity. Note also that $S \not\rightarrow W$ and $S \rightarrow Z$ implies $Z \not\rightarrow W$. But $Z \rightarrow A$ is assumed folded. \nmid . That is, $S \rightarrow A$ must be folded. Then we have both $S \rightarrow A$ and $Z \rightarrow A$ folded, though $S \neq Z$. That is, the attribute folding of A is not unique, even though we know by Lemma 2 that it must be unique. \nmid . Thus Σ^{q} must be left-reduced, altogether, therefore, parsimonious.

\square

A.2.5 Proof of Theorem 4

“Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, φ a total causal mapping over \mathcal{S} and Σ an fd set encoded through φ given \mathcal{S} . Then $x_a, x_b \in \mathcal{V}$ are such that x_b is causally dependent on x_a , i.e., $(x_a, x_b) \in C_\varphi^+$ iff there is some non-trivial fd $\langle X, B \rangle \in \Sigma^\triangleright$ with $A \in X$, where $B \mapsto x_b$ and $A \mapsto x_a$.”

Proof 26 We prove the statement by induction. We consider first the ‘if’ direction, and then its ‘only if’ converse.

(Base case). Let $\langle X, B \rangle \in \Sigma$ be some fd with $A \in X$, where $B \mapsto x_b$ and $A \mapsto x_a$. By Theorem 2, it is non-trivial and then by default (i.e., an empty application of R5) it is in Σ^\triangleright . But as $X \rightarrow B$ is in Σ , (Alg. 3) **h-encode** ensures that there is exactly one equation $f \in \mathcal{E}$ such that $\varphi(f) = x_b$ and $x_a \in \text{Vars}(f)$ where $B \mapsto x_b$ and $A \mapsto x_a$. Then by force of Eq. 3.1 we must have $(x_a, x_b) \in C_\varphi$. Thus, we obviously have $(x_a, x_b) \in C_\varphi^+$ as well.

(Induction). Now, recall Armstrong’s (R5) pseudo-transitivity rule adapted here for the particular case of singleton-rhs fd sets, viz., if $Y \rightarrow C$ and $CZ \rightarrow B$, then $YZ \rightarrow B$. By the inductive hypothesis, take any two non-trivial fd’s $\langle Y, C \rangle, \langle CZ, B \rangle \in \Sigma^\triangleright$ with $B \notin Y$ and assume that the causal dependency property holds for their attributes that encode variables. That is, let $D \in Y$ and $E \in Z$, where $D \mapsto x_d$ and $E \mapsto x_e$ for $x_d, x_e \in \mathcal{V}$ such that $(x_d, x_c), (x_e, x_b), (x_c, x_b) \in C_\varphi^+$. Note that both $Y \rightarrow C$ and $CZ \rightarrow B$ are non-trivial, then $C \notin Y$, $B \notin Z$ and $B \neq C$. Moreover, $B \notin Y$ has been assumed such that the fd $\langle YZ, B \rangle \in \Sigma^\triangleright$ to be derived by R5 over $Y \rightarrow C$ and $CZ \rightarrow B$ is also non-trivial to satisfy the condition of the theorem. Now, it is easy to see that the property holds likewise for non-trivial fd $\langle YZ, B \rangle \in \Sigma^\triangleright$. In fact, $(x_d, x_c), (x_c, x_b) \in C_\varphi^+$ implies $(x_d, x_b) \in C_\varphi^+$ and also by the inductive hypothesis we have $(x_e, x_b) \in C_\varphi^+$. That is, for either some $D \in Y$ or some $E \in Z$, we must have $(x_d, x_b), (x_e, x_b) \in C_\varphi^+$.

The converse ‘only if’ direction can be shown by a symmetrical inductive argument. That is, for the base case suppose $(x_a, x_b) \in C_\varphi$. Then, by Eq. 3.1 we know there is some $f \in \mathcal{E}$ such that $\varphi(f) = x_b$ and $x_a \in \text{Vars}(f)$. Moreover, in that case (Alg. 3) **h-encode** ensures there must be some non-trivial fd $\langle X, B \rangle \in \Sigma$ with $A \in X$ where $B \mapsto x_b$ and $A \mapsto x_a$. Thus by an empty application of R5 we

have $\langle X, B \rangle \in \Sigma^\triangleright$. The inductive step shows the property still holds for arbitrary causal dependencies in C_φ^+ . \square

A.2.6 Proof of Proposition 4

“Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a structure with variable $x \in \mathcal{V}$. Then x can only be a first cause of some $y \in \mathcal{V}$ if x is exogenous. Accordingly, any variable $y \in \mathcal{V}$ can only have some first cause $x \in \mathcal{V}$ if it is endogenous.”

Proof 27 The proof is straightforward from definitions. For the first statement, suppose by contradiction that $x \in \mathcal{V}$ is not exogenous but is a first cause of some $y \in \mathcal{V}$. By Def. 5, φ is bijective then there is some $f \in \mathcal{E}$ such that $\varphi(f) = x$. Moreover, as x is not exogenous then by Def. 8 it must be endogenous. In other words, there must be some $x_a \in \mathcal{V}$ such that $x_a \neq x \in \text{Vars}(f)$ and then by Eq. 3.1 we have $(x_a, x) \in C_\varphi$ hence $(x_a, x) \in C_\varphi^+$. However, as x is a first cause, by Def. 14 there can be no $y \in \mathcal{V}$ such that $(y, x) \in C_\varphi^+$. \nmid

Now, a symmetrical argument proves the second statement. Also by contradiction take a variable $y \in \mathcal{V}$ that is not endogenous and suppose it has some first cause $x \in \mathcal{V}$. As variable y is not endogenous then by Def. 8 it must be exogenous. In other words, there must be $f \in \mathcal{E}$ such that $\text{Vars}(f) = \{y\}$. Thus for any total causal mapping φ over \mathcal{S} , we must have $\varphi(f) = y$ and, for no $x \in \mathcal{V}$, we have $(x, y) \in C_\varphi$. Therefore it is not possible to derive $(x, y) \in C_\varphi^+$ for some $x \in \mathcal{V}$. But as y has some first cause $x \in \mathcal{V}$ by assumption, we must have $(x, y) \in C_\varphi^+$. \nmid \square

A.2.7 Proof of Lemma 3

“Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, φ a total causal mapping over \mathcal{S} and Σ an fd set encoded through φ given \mathcal{S} . Then a variable $x_a \in \mathcal{V}$ can only be a first cause of some variable $x_b \in \mathcal{V}$, where $\langle X, B \rangle \in \Sigma$, and $B \mapsto x_b$, $A \mapsto x_a$, if either (i) $A \in X$ or (ii) $A \notin X$ but there is $\langle Z, C \rangle \in \Sigma^\triangleright$ with $A \in Z$ and $C \in X$.”

Proof 28 We prove the statement by construction out of Theorem 4.

By Def. 14, one of the conditions for x_a to be a first cause of x_b is that $(x_a, x_b) \in C_\varphi^+$. Moreover, by Theorem 4 we know that $(x_a, x_b) \in C_\varphi^+$ can only hold if there is some non-trivial fd $\langle Z, B \rangle \in \Sigma^\triangleright$ with $A \in Z$, where $B \mapsto x_b$ and $A \mapsto x_a$. Now, by Def. 5 φ is bijective then there is $\langle X, B \rangle \in \Sigma$. Moreover, since

Σ is parsimonious, $X \rightarrow B$ is the only fd in Σ with B in its rhs. So let $A \notin X$. Then we know $X \neq Z$ hence $X \rightarrow B$ cannot be the fd required by Theorem 4.

Then such fd $Z \xrightarrow{b} B$ with $A \in Z$ can only exist if derived by some finite application of R5. That is, there must be some $\langle Z, C \rangle \in \Sigma$ with $A \in Z$ such that $X = CW$ for some W and then R5 can be applied over $\langle Z, C \rangle, \langle CW, B \rangle \in \Sigma$ to get non-trivial $\langle ZW, B \rangle \in \Sigma^\triangleright$ where $A \in Z$.

Now, it is easy to see that when such fd $Z \rightarrow C$ with $A \in Z$ does not exist in Σ^\triangleright (the second condition of the lemma), then obviously $Z \rightarrow C$ cannot exist in Σ to satisfy the requirement imposed by Theorem 4. That is, $(x_a, x_b) \notin C_\varphi^+$. \square

A.2.8 Proof of Theorem 5

“Let $\mathcal{S}(\mathcal{E}, \mathcal{V})$ be a complete structure, φ a total causal mapping over \mathcal{S} and Σ an fd set encoded through φ given \mathcal{S} . Now, let B be an attribute that encodes some variable $x_b \in \mathcal{V}$. If $\langle X, B \rangle \in \Upsilon(\Sigma)^{\triangleright, 1}$ then every first cause x_a of x_b (if any) is encoded by some attribute $A \in X$.”

Proof 29 We show that the existence of a missing first cause x_c of x_b for folded $X \xrightarrow{\triangleright} B$, where $B \mapsto x_b$ and $C \mapsto x_c$ but $C \notin X$ leads to a contradiction.

Suppose, by contradiction, that there is some missing first cause $x_c \in \mathcal{V}$ of x_b , where $C \mapsto x_c$ and $C \notin X$. Then, by Lemma 3, since variable x_c is a first cause of variable x_b , it must be exogenous and, for $\langle Y, B \rangle \in \Upsilon(\Sigma)$ either (i) $C \in Y$ or (ii) $C \notin Y$ but there is $\langle Z, D \rangle \in \Upsilon(\Sigma)^\triangleright$ with $C \in Z$ and some $D \in Y$.

In the first case (i), since x_c is exogenous and Σ is parsimonious, we have $\langle \phi, C \rangle \in \Sigma$ but by Def. 15 there can be no $W \rightarrow C$ in the v -projection $\Upsilon(\Sigma)$ of Σ . That is, C cannot be ‘consumed’ by R5 and then $\langle Y, B \rangle \in \Upsilon(\Sigma)$ with $C \in Y$ implies that, for any $\langle W, B \rangle \in \Upsilon(\Sigma)^\triangleright$, we must have $C \in W$. However, by assumption we have $\langle X, B \rangle \in \Upsilon(\Sigma)^{\triangleright}$ then, by Def. 12, $\langle X, B \rangle \in \Upsilon(\Sigma)^\triangleright$, yet $C \notin X$. \nmid

In the second case (ii), observe that $C \in Z$ and $D \in Y$, and let $Y = DS$. Then by R5 over $Z \rightarrow D$ and $DS \rightarrow B$ we get $\langle ZS, B \rangle \in \Upsilon(\Sigma)^\triangleright$, where $C \in Z$. Well, either $ZS \rightarrow B$ is folded or it is not, rendering two cases for analysis. If $ZS \rightarrow B$ is folded, then both $ZS \rightarrow B$ and $X \rightarrow B$ are folded. But as $\Upsilon(\Sigma)$ is

¹ Note that the folding is taken w.r.t. the v -projection of Σ , then x_b where $B \mapsto x_b$ is an endogenous variable.

parsimonious, then by Lemma 2 the folding of B must be unique. Therefore we must have $ZS = X$, with $C \in Z$ but $C \notin X$. \nmid .

Else, assume $ZS \rightarrow B$ is not folded. Then by Def. 11 there is some W with $W \not\subseteq ZS$ such that non-trivial $W \rightarrow ZS$ is in $\Upsilon(\Sigma)^+$ and $ZS \not\rightarrow W$.

However, as $C \in Z$ and $W \rightarrow ZS$, by (R3) decomposition we must have $W \rightarrow C$ in $\Upsilon(\Sigma)^+$, either with $C \in W$ or with $C \notin W$ and then $W \rightarrow C$ is non-trivial. But we know the latter cannot be the case by the same argument used in the first case (i), viz., x_c is exogenous with $C \mapsto x_c$ and Σ is parsimonious. That is, we must have $C \in W$. Furthermore, as $W \rightarrow ZS$ in $\Upsilon(\Sigma)^+$, then by R5 over $ZS \rightarrow B$ we get $\langle W, B \rangle \in \Upsilon(\Sigma)^\triangleright$ with $C \in W$. Now it is easy to see that the same situation recurs to $W \rightarrow B$. If it is not folded, eventually for some T we will have $T \xrightarrow{\triangleright} W \xrightarrow{\triangleright} B$ with $C \in T$, where $T \rightarrow B$ will be folded just like $X \rightarrow B$. That is, by (Lemma 2) the uniqueness of the folding of B , we will have $T = X$ with $C \in T$ and $C \notin X$. \nmid . \square

A.3. Proofs of Probabilistic DB Synthesis

A.3.1 Proof of Theorem 6

“Let \mathcal{S}_k and H_k be (resp.) the complete structure and ‘big’ fact table of hypothesis k , and let Γ'_k be the repaired factorization of \mathcal{S}_k over H_k , and Y_0 the ‘explanation’ table where hypothesis k is recorded. Now, let \mathbf{Y}_k be a U -relational schema defined $\mathbf{Y}_k \triangleq \text{synthesize4u}(\mathcal{S}_k, H_k, Y_0)$. Then \mathbf{Y}_k is in BCNF w.r.t. Γ'_k and is minimal-cardinality.”

Proof 30 Let $Y_k^i [V_i D_i \mid \phi A_i G_i]$ and $Y_k^j [\overline{V_j D_j} \mid S T]$ be (resp.) any u-factor projection and predictive projection of H_k . Note that all fd's in Γ'_k are either in $\Phi(\Gamma'_k)$ of form $\phi A_i \rightarrow B$ or $\phi \rightarrow B$, or in $\Upsilon(\Gamma'_k)$ of form $A_1 A_2 \dots A_\ell S \rightarrow T$ with $v \in S$. We must show that no fd in $(\Gamma'_k)^+$ can violate Y_k^i or Y_k^j . It is easy to see that the projection (cf. Def. 20) of non-trivial fd's in $\Phi(\Gamma'_k)^+$ onto Y_k^j is empty, just like the projection of non-trivial fd's in $\Upsilon(\Gamma'_k)^+$ onto Y_k^i .

For the u-factor projections, note by Def. 21 that for any fd $X \rightarrow C$ in $(\Gamma'_k)^+$ to violate BCNF in $Y_k^i [V_i D_i \mid \phi A_i G_i]$, it must be non-trivial ($C \notin X$) with $XC \subseteq \phi A_i G_i$ but $X \not\rightarrow \phi A_i G_i$ (that is, X is not a superkey for Y_k^i). Note that we

have both $\phi A_i \rightarrow B$ and $\phi \rightarrow A_i B$ in $(\Gamma'_k)^+$ for any $B \in G_i$, but both ϕA_i and ϕ are superkeys for Y_k^i . Also, that there can be no non-trivial fd's $\langle X, C \rangle \in \Phi(\Gamma'_k)^+$ with $\phi \not\subseteq X$, and by the definition of Problem 1 we know that $A_i G_i$ is a maximal group. So, for any non-trivial $\langle X, C \rangle \in \Phi(\Gamma'_k)^+$, X must be a superkey for Y_k^i . Thus no u-factor projection can be subject of BCNF violation w.r.t. Γ'_k .

Now, for predictive projection $Y_k^j [\overline{V_j D_j} \mid S T]$ let us reconstruct the process towards deriving $\langle S, T \rangle \in \Upsilon(\Gamma'_k)^\triangleright$. Note that it is derived by **synthesize4u** simulating ℓ applications of R5 over $\langle \phi, A_i \rangle$, $\langle A_1 A_2 \dots A_\ell S, T \rangle \in \Gamma'_k$ for $0 \leq i \leq \ell$. Note also that (i) no cyclic fd's can be involved in such R5 applications, as they must always be over an fd in $\Phi(\Gamma'_k)$; and (ii) as a result of (Alg. 7) **merge**, $A_1 A_2 \dots A_\ell S \rightarrow T$ was the only non-trivial fd in the projection $\pi_{A_1 A_2 \dots A_\ell S T}(\Gamma'_k)$. Thus, the only non-trivial fd's in the projection $\pi_{ST}((\Gamma'_k))^+$ in addition to $S \rightarrow T$ itself must be of form $S \rightarrow C$ rendered out of (R3) decomposition from it for all $C \in T$. In any such fd's, we have S as a superkey for Y_k^j . Therefore no predictive projection can be subject of BCNF violation w.r.t. Γ'_k .

For the minimality note, as a consequence of (Alg. 7) **merge**, any two schemes $Y_k^p[XZ]$, $Y_k^q[VW]$ are rendered by **synthesize4u** into \mathbf{Y}_k iff we have fd's $\langle X, Z \rangle, \langle V, W \rangle \in (\Gamma'_k)^+$ and $X \not\rightarrow V$, i.e., it is not the case that both $X \rightarrow V$ and $V \rightarrow X$ hold in $(\Gamma'_k)^+$. Now, to prove that \mathbf{Y}_k is minimal-cardinality, we have to find that merging any such pair of arbitrary schemes shall hinder BCNF in \mathbf{Y}_k . In fact, take $\mathbf{Y}'_k := \mathbf{Y}_k \setminus (Y_k^p[XZ] \cup Y_k^q[VW]) \cup Y_k^\ell[XZVW]$. As $X \not\rightarrow V$, then neither X nor V can be a superkey for Y_k^ℓ , which therefore cannot be in BCNF. \square

A.3.2 Proof of Theorem 7

“Let \mathcal{S}_k be the complete structure of hypothesis k , and $H_k[U]$ its ‘big’ fact table such that Γ'_k is the repaired factorization of \mathcal{S}_k over H_k and Y_0 is the ‘explanation’ table where hypothesis k is recorded. Now, let \mathbf{Y}_k be a U -relational schema defined $\mathbf{Y}_k \triangleq \text{synthesize4u}(\mathcal{S}_k, H_k, Y_0)$. Then,

- (a) *the join $\bowtie_{i=1}^m Y_k^i [V_i D_i \mid \phi A_i G_i]$ of any subset of the u-factor projections of H_k is lossless w.r.t. Γ'_k .*
- (b) *any predictive projection $Y_k^j [\overline{V_j D_j} \mid S T]$, result of a join of the theoretical u-factor $Y_0 [V_0 D_0 \mid \phi v]$ with the ‘big’ fact table $H_k[U]$ and in turn with*

u-factor projections $Y_k^i[\overline{V_i D_i} | \phi A_i G_i]$, is lossless w.r.t. Γ'_k .

Proof 31 For item (a), by Lemma 6, we know that any pair $Y_k^i[V_i D_i | \phi A_i G_i]$, $Y_k^j[V_j D_j | \phi A_j G_j]$ of u-factor projections of H_k will have a lossless join w.r.t. Γ'_k iff $(\phi A_i G_i \cap \phi A_j G_j) \rightarrow (\phi A_i G_i \setminus \phi A_j G_j)$ or $(\phi A_i G_i \cap \phi A_j G_j) \rightarrow (\phi A_j G_j \setminus \phi A_i G_i)$ hold in $(\Gamma'_k)^+$. By Def. 17, we know that $(\phi A_i G_i \cap \phi A_j G_j) = \{\phi\}$, and $\phi A_i G_i \setminus \phi A_j G_j = A_i G_i$. In fact $\phi \rightarrow A_i G_i$ is a repaired fd in Γ'_k , therefore Y_k^i and Y_k^j have a lossless join. Now, since the join is an associative operation [20, p. 62], and as we have chosen Y_k^i and Y_k^j arbitrarily, then clearly any subset of the u-factor projections must have a lossless join.

For item (b), for any predictive projection $Y_k^j[\overline{V_j D_j} | ST]$ take the join $\bowtie Y_k^i[\overline{V_i D_i} | \phi \overline{A_i G_i}]$ of u-factor projections such that, for all $A_i \in \overline{A_i G_i}$, we have $A_i \in W \subset Z$ where $S = Z \setminus W$ and $\langle Z, T \rangle \in \Gamma'_k$. That is, A_i is a pivot attribute representing a first cause of some $C \in T$. By item (a), we know that such join is lossless.

We must show that the join $\bowtie Y_k^i$ with ‘big’ fact table $H_k[U]$ is also lossless. By Lemma 6, that is the case iff $(\phi \overline{A_i G_i} \cap U) \rightarrow (\phi \overline{A_i G_i} \setminus U)$ or $(\phi \overline{A_i G_i} \cap U) \rightarrow (U \setminus \phi \overline{A_i G_i})$ hold in $(\Gamma'_k)^+$. In fact, we have $(\phi \overline{A_i G_i} \cap U) = \phi \overline{A_i G_i}$ and $(\phi \overline{A_i G_i} \setminus U) = \emptyset$ such that $\phi \overline{A_i G_i} \rightarrow \emptyset$ is trivially in $(\Gamma'_k)^+$.

Finally, the join of theoretical u-factor $Y_0[V_0 D_0 | \phi v]$ with big fact table $H_k[U]$ must be lossless likewise. In fact, note that $(\phi v \cap U) = \phi v$, and $(\phi v \setminus U) = \emptyset$. Then also trivially we have $\phi v \rightarrow \emptyset$, which is in $(\Gamma'_k)^+$ as well. Since the join is commutative [20, p. 62], the order of application is irrelevant therefore the join of all joins examined above taken together must be lossless. \square

Lemma 6 Let Σ be a set of fd’s on attributes U , and $R_i[S], R_j[T] \in \mathbf{R}[U]$ be relation schemes with $ST \subseteq U$; and let $\pi_{ST}(\Sigma)$ be the projection of Σ onto ST . Then $R_i[S]$ and $R_j[T]$ have a lossless join w.r.t. $\pi_{ST}(\Sigma)$ iff $(S \cap T) \rightarrow (S \setminus T)$ or $(S \cap T) \rightarrow (T \setminus S)$ hold in $\pi_{ST}(\Sigma)^+$.

Proof 32 See Ullman [20, p. 397]. \square